

# Sample Complexity of Learning Mahalanobis Distance Metrics

Nakul Verma\* and Kristin Branson†

Janelia Research Campus  
Howard Hughes Medical Institute, Virginia, USA

## Abstract

Metric learning seeks a transformation of the feature space that enhances prediction quality for the given task at hand. In this work we provide PAC-style sample complexity rates for supervised metric learning. We give matching lower- and upper-bounds showing that the sample complexity scales with the representation dimension when no assumptions are made about the underlying data distribution. However, by leveraging the structure of the data distribution, we show that one can achieve rates that are *fine-tuned* to a specific notion of intrinsic complexity for a given dataset. Our analysis reveals that augmenting the metric learning optimization criterion with a simple norm-based regularization can help adapt to a dataset’s intrinsic complexity, yielding better generalization. Experiments on benchmark datasets validate our analysis and show that regularizing the metric can help discern the signal even when the data contains high amounts of noise.

## 1 Introduction

In many machine learning tasks, data is represented in a high-dimensional Euclidean space where each dimension corresponds to some interesting measurement of the observation. Often, practitioners include a variety of measurements in hopes that some combination of these features will capture the relevant information. While it is natural to represent such data in a Real space of measurements, there is no reason to expect that using Euclidean ( $L_2$ ) distances to compare the observations will be necessarily useful for the task at hand. Indeed, the presence of uninformative or mutually correlated measurements simply inflates the  $L_2$ -distances between pairs of observations, rendering distance-based comparisons ineffective.

*Metric learning* has emerged as a powerful technique to learn a good notion of distance or a *metric* in the representation space that can emphasize the feature combinations that help in the prediction task while suppressing the contribution of spurious measurements. The past decade has seen a variety of successful metric learning algorithms that leverage various attributes of the problem domain. A few notable examples include exploiting class labels to find a Mahalanobis distance metric that maximizes the distance between dissimilar observations while minimizing distances between similar ones to improve classification quality (Weinberger & Saul, 2009; Davis et al., 2007),

---

\*email: verman@janelia.hhmi.org; corresponding author.

†email: bransonk@janelia.hhmi.org

and explicitly optimizing for a downstream prediction task such as information retrieval (McFee & Lanckriet, 2010).

Despite the popularity of metric learning methods, few studies have focused on studying how the problem complexity scales with key attributes of a given dataset. For instance, how do we expect the generalization error to scale—both theoretically and practically—as one varies the number of informative and uninformative measurements, or changes the noise levels?

Here we study supervised metric learning more formally and gain a better understanding of how different modalities in data affect the metric learning problem. We develop two general frameworks for PAC-style analysis of supervised metric learning. We can categorize the popular metric learning algorithms into an empirical error minimization problem in one of the two frameworks. The first generic framework, the distance-based metric learning framework, uses class label information to derive distance constraints. The key objective is to learn a metric that on average yields smaller distances between examples from the same class than those from different classes. Some popular algorithms that optimize for such distance-based objectives include Mahalanobis Metric for Clustering (MMC) by Xing et al. (2002) and Information Theoretic Metric Learning (ITML) by Davis et al. (2007). Instead of using distance comparisons as a proxy, however, one can also optimize for a specific prediction task directly. The second generic framework, the classifier-based metric learning framework, explicitly incorporates the hypothesis associated with the prediction task of interest to learn effective distance metrics. A few interesting examples in this regime include the work by McFee & Lanckriet (2010) that finds metrics that improve ranking quality in information retrieval tasks, and the work by Shaw et al. (2011) that learns metrics that help predict connectivity structure in networked data.

Our analysis shows that in both frameworks, the sample complexity scales with the representation dimension for a given dataset (Lemmas 1 and 3), and this dependence is necessary in the absence of any specific assumptions on the underlying data distribution (Lemmas 2 and 4). By considering any Lipschitz loss, our results generalize previous sample complexity results (see our discussion in Section 6) and, for the first time in the literature, provide matching lower bounds.

In light of the observation made earlier that data measurements often include uninformative or weakly informative features, we expect a metric that yields good generalization performance to de-emphasize such features and accentuate the relevant ones. We can thus formalize the *metric learning complexity* of a given dataset in terms of the intrinsic complexity  $d$  of the metric that reweights the features in a way that yields the best generalization performance. (For Mahalanobis distance metrics, we can characterize the intrinsic complexity by the *norm* of the matrix representation of the metric.) We refine our sample complexity result and show a *dataset-dependent* bound for both frameworks that scales with dataset’s intrinsic metric learning complexity  $d$  (Corollary 7).

Taking guidance from our dataset-dependent result, we propose a simple variation on the empirical risk minimizing (ERM) algorithm that, when given an i.i.d. sample, returns a metric (of complexity  $\hat{d}$ ) that jointly minimizes the observed sample bias and the expected intra-class variance for metrics of fixed complexity  $\hat{d}$ . This bias-variance balancing algorithm can be viewed as a structural risk minimizing algorithm that provides better generalization performance than an ERM algorithm and justifies norm-regularization of weighting metrics in the optimization criteria for metric learning.

Finally, we evaluate the practical efficacy of our proposed norm-regularization criteria with some popular metric learning algorithms on benchmark datasets (Section 5). Our experiments highlight that the norm-regularization indeed helps in learning weighting metrics that better adapt to the signal in data in high-noise regimes.

## 2 Preliminaries

Given a representation space  $X = \mathbb{R}^D$  of  $D$  real-valued measurements of observations of interest, the goal of metric learning is to learn a *metric*  $M$  (that is, a  $D \times D$  real-valued weighting matrix on  $X$ ; to remove arbitrary scaling we shall assume that the maximum singular value of  $M$ , that is,  $\sigma_{\max}(M) = 1$ )<sup>1</sup> that minimizes some notion of *error* on data drawn from an unknown underlying distribution  $\mathcal{D}$  on  $X \times \{0, 1\}$ . Specifically, we want to find the metric

$$M^* := \operatorname{argmin}_{M \in \mathcal{M}} \operatorname{err}(M, \mathcal{D}),$$

from the class of metrics  $\mathcal{M}$  under consideration, that is,  $\mathcal{M} := \{M \mid M \in \mathbb{R}^{D \times D}, \sigma_{\max}(M) = 1\}$ . For supervised metric learning, this *error* is typically label-based and can be defined in multiple reasonable ways. As discussed earlier, we explore two intuitive regimes for defining error.

**Distance-based error.** A popular criterion for quantifying error in metric learning is by comparing *distances* amongst points drawn from the underlying data distribution. Ideally, we want a weighting metric  $M$  that brings data from the same class closer together than those from opposite classes. In a distance-based framework, a natural way to accomplish this is to find a weighting  $M$  that yields shorter distances between pairs of observations from the same class than those from different classes. By penalizing how often and by how much the distances violate these constraints gives rise to the particular form of the error.

Let the variable  $z = (x, y)$  denote a random draw from  $\mathcal{D}$  with  $x \in X$  as the observation and  $y \in \{0, 1\}$  its associated label, and let  $\lambda$  denote how severely one wants to penalize the distance violations, then a natural definition of *distance-based* error becomes:

$$\operatorname{err}_{\text{dist}}^\lambda(M, \mathcal{D}) := \mathbb{E}_{z_1, z_2 \sim \mathcal{D}} \left[ \phi^\lambda(\rho_M(x_1, x_2), Y) \right],$$

for a generic distance-based loss function  $\phi^\lambda(\rho_M, Y)$ , that computes the degree of violation between weighted distance  $\rho_M(x_1, x_2) := \|M(x_1 - x_2)\|^2$  and the label agreement  $Y := \mathbf{1}[y_1 = y_2]$  among a pair  $z_1 = (x_1, y_1)$  and  $z_2 = (x_2, y_2)$  drawn from  $\mathcal{D}$ .

An example instantiation of  $\phi$  popular in literature encourages metrics that yield distances that are no more than some upper limit  $U$  between observations from the same class, and distances that are no less than some lower limit  $L$  between those from different classes (for some  $U < L$ ). Thus

$$\phi_{L,U}^\lambda(\rho_M, Y) := \begin{cases} \min\{1, \lambda[\rho_M - U]_+\} & \text{if } Y = 1 \\ \min\{1, \lambda[L - \rho_M]_+\} & \text{otherwise} \end{cases}, \quad (1)$$

where  $[A]_+ := \max\{0, A\}$ .

Xing et al. (2002) optimize an efficiently computable variant of this criterion, in which they look for a metric that keeps the total pairwise distance amongst the observations from the same class less than a constant while maximizing the total pairwise distance amongst the observations from opposite classes. The variant proposed by Davis et al. (2007) explicitly includes the upper and lower limits with an added regularization on the learned  $M$  to be close to a pre-specified metric of interest  $M_0$ .

While we discuss loss-functions  $\phi$  that handle distances between a *pair* of observations, it is easy to extend to distances among *triplets*. Rather than having hard upper and lower limits which every

<sup>1</sup>Note that we are looking at the linear form of the metric  $M$ ; usually the corresponding quadratic form  $M^\top M$  is discussed in the literature, which is necessarily positive semi-definite.

pair of the same and the opposite classes must obey, a triplet-based comparison typically focuses on relative distances between three observations at a time. A natural instantiation in this case becomes:

$$\phi_{\text{triplet}}^{\lambda}(\rho_M(x_1, x_2), \rho_M(x_1, x_3), (y_1, y_2, y_3)) := \begin{cases} \min\{1, \lambda[\rho_M(x_1, x_2) - \rho_M(x_1, x_3)]_+\} & \text{if } y_1 = y_2 \neq y_3 \\ 0 & \text{otherwise} \end{cases},$$

for a triplet  $(x_1, y_1), (x_2, y_2), (x_3, y_3)$  drawn from  $\mathcal{D}$ .

Weinberger & Saul (2009) discuss an interesting variant of this, in which instead of looking at all triplets in a given training sample, they focus on triplets of observations in local neighborhoods and learn a metric that maintains a gap or a margin among distances between observations from the same class and those from the opposite class. Improving the quality of distance comparisons in local neighborhoods directly affects the nearest neighbor performance, making this a popular technique.

**Classifier-based Error.** Distance comparisons typically act as a surrogate for a specific downstream prediction task. If we want a metric that directly optimizes for a task, we need to explicitly incorporate the hypothesis class being used for that task while finding a good weighting metric.

This simple but effective insight has been used recently by McFee & Lanckriet (2010) for improving ranking results in information retrieval problems by explicitly incorporating ranking losses while learning an effective weighting metric. Shaw et al. (2011) also follow this principle and explicitly include network topology constraints to learn a weighting metric that can better predict the connectivity structure in social and web networks.

We can formalize the classifier-based metric learning framework by considering a fixed hypothesis class  $\mathcal{H}$  of interest on the measurement domain. To keep the discussion general, we shall assume that the hypotheses are real-valued and can be regarded as a measure of confidence in classification, that is, each  $h \in \mathcal{H}$  is of the form  $h : X \rightarrow [0, 1]$ . (One can obtain the binary predictions from  $h$  by a simple thresholding at  $1/2$ .) Then, the error induced by a particular weighting metric  $M$  on the measurement space  $X$  can be defined as the best possible error that can be obtained by hypotheses in  $\mathcal{H}$ , that is

$$\text{err}_{\text{hypoth}}(M, \mathcal{D}) := \inf_{h \in \mathcal{H}} \mathbb{E}_{(x, y) \sim \mathcal{D}} \left[ \mathbf{1} [|h(Mx) - y| \geq 1/2] \right].$$

We shall study how this error scales with various key parameters of the metric learning problem.

### 3 Learning a Metric from Samples

In any practical setting, we estimate the ideal weighting metric  $M^*$  by minimizing the empirical version of the error criterion from a finite size sample from  $\mathcal{D}$ .

Let  $S_m$  denote a sample of size  $m$ , and  $\text{err}(M, S_m)$  denote the empirical error on the sample  $S_m$  (the exact definitions of  $S_m$  and the form of  $\text{err}(M, S_m)$  are discussed later). We can then define the empirical risk minimizing metric based on  $m$  samples as  $M_m^* := \text{argmin}_M \text{err}(M, S_m)$ . Most practical algorithms, of course, return some approximation of  $M_m^*$ , and thus it is important to compare the generalization ability of  $M_m^*$  to that of theoretically optimal  $M^*$ . That is, how

$$\text{err}(M_m^*, \mathcal{D}) - \text{err}(M^*, \mathcal{D}) \tag{2}$$

converges as the sample size  $m$  grows.

### 3.1 Distance-Based Error Analysis

Given an i.i.d. sequence of observations  $z_1, z_2, \dots$  from  $\mathcal{D}$ , we can pair the observations together to form a *paired* sample  $S_m = \{(z_1, z_2), (z_3, z_4), \dots, (z_{2m-1}, z_{2m})\} = \{(z_{1,i}, z_{2,i})\}_{i=1}^m$  of size  $m$ , and define the sample based distance error  $\text{err}_{\text{dist}}^\lambda(M, S_m)$  induced by a metric  $M$  as

$$\text{err}_{\text{dist}}^\lambda(M, S_m) := \frac{1}{m} \sum_{i=1}^m \phi^\lambda(\rho_M(x_{1,i}, x_{2,i}), \mathbf{1}[y_{1,i} = y_{2,i}]).$$

Then for any bounded support distribution  $\mathcal{D}$  (that is, each  $(x, y) \sim \mathcal{D}$ ,  $\|x\| \leq B < \infty$ ), we have the following convergence result.<sup>2</sup>

**Lemma 1** *Fix any sample size  $m$ , and let  $S_m$  be an i.i.d. paired sample of size  $m$  from an unknown bounded distribution  $\mathcal{D}$  (with bound  $B$ ). For any distance-based loss function  $\phi^\lambda$  that is  $\lambda$ -Lipschitz in the first argument, with probability at least  $1 - \delta$  over the draw of  $S_m$ ,*

$$\sup_{M \in \mathcal{M}} [\text{err}_{\text{dist}}^\lambda(M, \mathcal{D}) - \text{err}_{\text{dist}}^\lambda(M, S_m)] \leq O\left(\lambda B^2 \sqrt{\frac{D \ln(1/\delta)}{m}}\right).$$

Using this lemma we can get the desired convergence rate (Eq. 2). Fix  $M^* \in \mathcal{M}$ , then for any  $0 < \delta < 1$  and  $m \geq 1$ , with probability at least  $1 - \delta$ , we have

$$\begin{aligned} \text{err}_{\text{dist}}^\lambda(M_m^*, \mathcal{D}) - \text{err}_{\text{dist}}^\lambda(M^*, \mathcal{D}) &= \text{err}_{\text{dist}}^\lambda(M_m^*, \mathcal{D}) - \text{err}_{\text{dist}}^\lambda(M_m^*, S_m) + \text{err}_{\text{dist}}^\lambda(M_m^*, S_m) - \text{err}_{\text{dist}}^\lambda(M^*, S_m) \\ &\quad + \text{err}_{\text{dist}}^\lambda(M^*, S_m) - \text{err}_{\text{dist}}^\lambda(M^*, \mathcal{D}) \\ &\leq O\left(\lambda B^2 \sqrt{\frac{D \ln(1/\delta)}{m}}\right) + \sqrt{\frac{\ln(2/\delta)}{2m}} \\ &= O\left(\lambda B^2 \sqrt{\frac{D \ln(1/\delta)}{m}}\right), \end{aligned}$$

by noting (i)  $\text{err}_{\text{dist}}^\lambda(M_m^*, S_m) \leq \text{err}_{\text{dist}}^\lambda(M^*, S_m)$ , since  $M_m^*$  is empirical error minimizing on  $S_m$ , and (ii) by using Hoeffding's inequality on the fixed  $M^*$  to conclude that with probability at least  $1 - \delta/2$ ,  $\text{err}_{\text{dist}}^\lambda(M^*, S_m) - \text{err}_{\text{dist}}^\lambda(M^*, \mathcal{D}) \leq \sqrt{\frac{\ln(2/\delta)}{2m}}$ .

Thus to achieve a specific estimation error rate  $\epsilon$ , the number of samples  $m = \Omega\left(\left(\frac{\lambda B^2}{\epsilon}\right)^2 D \ln\left(\frac{1}{\delta}\right)\right)$  are sufficient to conclude, with confidence at least  $1 - \delta$ , the empirical risk minimizing metric  $M_m^*$  will have estimation error of at most  $\epsilon$ . This shows that one never needs more than a number proportional to the representation dimension  $D$  examples to achieve the desired level of accuracy.

Since typical applications have a large representation dimension, it is instructive to study if such a strong dependency on  $D$  necessary. It turns out that even for simple distance-based loss functions like  $\phi_{L,U}^\lambda$  (cf. Eq. 1), there are data distributions for which one cannot get away with fewer than linear in  $D$  samples and ensure good estimation errors. In particular we have the following.

**Lemma 2** *Let  $\mathcal{A}$  be any algorithm that, given an i.i.d. sample  $S_m$  (of size  $m$ ) from a fixed unknown bounded support distribution  $\mathcal{D}$ , returns a weighting metric from  $\mathcal{M}$  that minimizes the empirical*

<sup>2</sup>We only present the results for paired distance comparisons; the results are easily extended to triplet-based comparisons.

error with respect to distance-based loss function  $\phi_{L,U}^\lambda$ . There exist  $\lambda \geq 0$ ,  $0 \leq U < L$ , such that for all  $0 < \epsilon, \delta < 1/64$ , there exists a bounded support distribution  $\mathcal{D}$ , such that if  $m \leq \frac{D+1}{512\epsilon^2}$  then

$$\mathbf{P}_{S_m} \left[ \text{err}_{\text{dist}}^\lambda(\mathcal{A}(S_m), \mathcal{D}) - \text{err}_{\text{dist}}^\lambda(M^*, \mathcal{D}) > \epsilon \right] > \delta.$$

While this may seem discouraging for large-scale applications of metric learning, note that here we made no assumptions about the underlying structure of the data distribution  $\mathcal{D}$ , making this a worst-case analysis. As the individual features in real-world datasets contain varying amounts of information for good classification performance, one hopes for a more relaxed dependence on  $D$  for metric learning in these settings. This is explored in Section 4.

### 3.2 Classifier-Based Error Analysis

In this setting, we can use an i.i.d. sequence of observations  $z_1, z_2, \dots$  from  $\mathcal{D}$  to obtain the sample  $S_m = \{z_i\}_{i=1}^m$  of size  $m$  directly. To analyze the generalization ability of the weighting metrics optimized with respect to an underlying hypothesis class  $\mathcal{H}$ , we need to effectively analyze the classification complexity of  $\mathcal{H}$ . The scale sensitive version of VC-dimension, also known as the “fat-shattering dimension”, of a real-valued hypothesis class (denoted by  $\text{Fat}_\gamma(\mathcal{H})$ ) encodes the right notion of classification complexity and provides an intuitive way to relate the generalization error to the empirical error at a margin  $\gamma$  (see for instance the work of Anthony & Bartlett (1999) for an excellent discussion).

In the context of metric learning with respect to a fixed hypothesis class, define the empirical error at a margin  $\gamma$  as

$$\text{err}_{\text{hypo}}^\gamma(M, S_m) := \inf_{h \in \mathcal{H}} \frac{1}{m} \sum_{(x_i, y_i) \in S_m} \mathbf{1}[\text{Margin}(h(Mx_i), y_i) < \gamma],$$

$$\text{where } \text{Margin}(\hat{y}, y) := \begin{cases} \hat{y} - 1/2 & \text{if } y = 1 \\ 1/2 - \hat{y} & \text{otherwise} \end{cases}.$$

Then for any bounded support distribution  $\mathcal{D}$  (that is, each  $(x, y) \sim \mathcal{D}$ ,  $\|x\| \leq B < \infty$ ), we have the following convergence result that relates the estimation error rate of the weighting metrics with that of the fat-shattering dimension of the underlying base hypothesis class.

**Lemma 3** *Let  $\mathcal{H}$  be a  $\lambda$ -Lipschitz base hypothesis class. Pick any  $0 < \gamma < 1/2$ , and let  $m \geq \text{Fat}_{\gamma/16}(\mathcal{H}) \geq 1$ . Then with probability at least  $1 - \delta$  over an i.i.d. draw of sample  $S_m$  (of size  $m$ ) from a bounded unknown distribution  $\mathcal{D}$  (with bound  $B$ ) on  $X \times \{0, 1\}$ ,*

$$\sup_{M \in \mathcal{M}} \left[ \text{err}_{\text{hypo}}(M, \mathcal{D}) - \text{err}_{\text{hypo}}^\gamma(M, S_m) \right] \leq O \left( \sqrt{\frac{1}{m} \ln \frac{1}{\delta} + \frac{D^2}{m} \ln \frac{D}{\epsilon_0} + \frac{\text{Fat}_{\gamma/16}(\mathcal{H})}{m} \ln \left( \frac{m}{\gamma} \right)} \right).$$

where  $\epsilon_0 := \min\{\frac{\gamma}{2}, \frac{1}{2\lambda B}\}$ , and  $\text{Fat}_{\gamma/16}(\mathcal{H})$  is the fat-shattering dimension of the base hypothesis class  $\mathcal{H}$  at margin  $\gamma/16$ .

Using a similar line of argument as before, we can bound the key quantity of interest (Eq. 2) and conclude for any  $0 < \gamma < 1/2$  and any  $m \geq 1$ , with probability  $\geq 1 - \delta$

$$\text{err}_{\text{hypo}}(M_m^*, \mathcal{D}) - \text{err}_{\text{hypo}}^\gamma(M^*, \mathcal{D}) = O \left( \sqrt{\frac{D^2 \ln(D/\epsilon_0)}{m} + \frac{\text{Fat}_{\gamma/16}(\mathcal{H}) \ln(m/\delta\gamma)}{m}} \right).$$

Here  $\epsilon_0 = \min\{\frac{\gamma}{2}, \frac{1}{2\lambda B}\}$  for a  $\lambda$ -Lipschitz hypothesis class  $\mathcal{H}$ . Thus to achieve a specific estimation error rate  $\epsilon$ , the number of samples  $m = \Omega\left(\frac{D^2 \ln(\lambda DB/\gamma) + \text{Fat}_{\gamma/16}(\mathcal{H}) \ln(1/\delta\gamma)}{\epsilon^2}\right)$  suffices to say, with confidence at least  $1 - \delta$ , the empirical risk minimizing metric  $M_m^*$  will have estimation error at most  $\epsilon$ .

It is interesting to note that the task of finding an optimal metric only additively increases the sample complexity over the complexity of finding the optimal hypothesis from the underlying hypothesis class.

In contrast to the sample complexity of distance-based framework (c.f. Lemma 1), here we get a quadratic dependence on the representation dimension. The following lemma shows that a strong dependence on the representation dimension is necessary in absence of any specific assumptions on the underlying data distribution and the base hypothesis class.

**Lemma 4** *Pick any  $0 < \gamma < 1/8$ . Let  $\mathcal{H}$  be a base hypothesis class of  $\lambda$ -Lipschitz functions mapping from  $X = \mathbb{R}^D$  into the interval  $[1/2 - 4\gamma, 1/2 + 4\gamma]$  that is closed under addition of constants. That is*

$$h \in \mathcal{H} \implies h' \in \mathcal{H}, \text{ where } h' : x \rightarrow h(x) + c \text{ for all } c.$$

*Then for any classification algorithm  $\mathcal{A}$ , and for any  $B \geq 1$ , there exists  $\lambda \geq 0$ , for all  $0 < \epsilon, \delta < 1/64$ , there exists a bounded support distribution  $\mathcal{D}$  (with bound  $B$ ) such that if  $m \ln^2 m < O\left(\frac{D^2 + d}{\epsilon^2 \ln(1/\gamma^2)}\right)$*

$$\mathbf{P}_{S_m \sim \mathcal{D}}[\text{err}_{\text{hypoth}}(h^*, \mathcal{D}) > \text{err}_{\text{hypoth}}^\gamma(\mathcal{A}(S_m), \mathcal{D}) + \epsilon] > \delta,$$

*where  $d := \text{Fat}_{768\gamma}(\mathcal{H})$  is the fat-shattering dimension of  $\mathcal{H}$  at margin  $768\gamma$ .*

## 4 Data with Uninformative and Weakly Informative Features

Different measurements have varying degrees of “information content” for the particular supervised classification task of interest. Any algorithm or analysis that studies the design of effective comparisons between observations must account for this variability.

To get a solid footing for our study, we introduce the concept of *metric learning complexity* of a given dataset. Our key observation is that a metric that yields good generalization performance should emphasize relevant features while suppressing the contribution of spurious features. Thus, a good metric reflects the quality of individual feature measurements of data and their relative value for the learning task. We can leverage this and define the metric learning complexity of a given dataset as the *intrinsic complexity*  $d$  of the weighting metric that yields the best generalization performance for that dataset (if multiple metrics yield best performance, we select the one with minimum  $d$ ). A natural way to characterize the intrinsic complexity of a weighting metric  $M$  is via the norm of the matrix representation of  $M$ . Using metric learning complexity as our gauge for the richness of the feature set in a given dataset, we can refine our analysis in both our canonical metric learning frameworks.

### 4.1 Distance-Based Refinement

We start with the following refinement of the distance-based metric learning sample complexity for a class of Frobenius norm-bounded weighting metrics.

**Lemma 5** Let  $\mathcal{M}$  be any class of weighting metrics on the feature space  $X = \mathbb{R}^D$ . Fix any sample size  $m$ , and let  $S_m$  be an i.i.d. paired sample of size  $m$  from an unknown bounded distribution  $\mathcal{D}$  on  $X \times \{0, 1\}$  (with bound  $B$ ). For any distance-based loss function  $\phi^\lambda$  that is  $\lambda$ -Lipschitz in the first argument, with probability at least  $1 - \delta$  over the draw of  $S_m$ ,

$$\sup_{M \in \mathcal{M}} [\text{err}_{\text{dist}}^\lambda(M, \mathcal{D}) - \text{err}_{\text{dist}}^\lambda(M, S_m)] \leq O\left(\lambda B^2 \sqrt{\frac{d \ln(1/\delta)}{m}}\right),$$

where  $d$  is a uniform upperbound on the Frobenius norm of the quadratic form of weighting metrics in  $\mathcal{M}$ , that is,  $\sup_{M \in \mathcal{M}} \|M^\top M\|_F^2 \leq d$ .

Observe that if our dataset has a low metric learning complexity (say,  $d \ll D$ ), then considering an appropriate class of norm-bounded weighting metrics can help sharpen the sample complexity result, yielding a *dataset-dependent* bound. We discuss how to automatically adapt to the right complexity class in Section 4.3 below.

## 4.2 Classifier-Based Refinement

Effective data-dependent analysis of classifier-based metric learning requires accounting for potentially complex interactions between an arbitrary base hypothesis class and the distortion induced by a weighting metric to the unknown underlying data distribution. To make the analysis tractable while still keeping our base hypothesis class  $\mathcal{H}$  general, we shall assume that  $\mathcal{H}$  is a class of two layer feed-forward neural networks. Recall that for any smooth target function  $f^*$ , a two layer feed-forward neural network (with appropriate number of hidden units and connection weights) can approximate  $f^*$  arbitrarily well (Hornik et al., 1989), so this class is flexible enough to incorporate most reasonable target hypotheses.

More formally, define the base hypothesis class of two layer feed-forward neural network with  $K$  hidden units as

$$\mathcal{H}_{\sigma^\gamma}^{2\text{-net}} := \left\{ x \mapsto \sum_{i=1}^K w_i \sigma^\gamma(v_i \cdot x) \mid \|w\|_1 \leq 1, \|v_i\|_1 \leq 1 \right\},$$

where  $\sigma^\gamma : \mathbb{R} \rightarrow [-1, 1]$  is a smooth, strictly monotonic,  $\gamma$ -Lipschitz activation function with  $\sigma^\gamma(0) = 0$ . Then for the generalization error of a weighting metric  $M$  defined with respect to any classifier-based  $\lambda$ -Lipschitz loss function  $\phi^\lambda$

$$\text{err}_{\text{hypo}}^\lambda(M, \mathcal{D}) := \inf_{h \in \mathcal{H}_{\sigma^\gamma}^{2\text{-net}}} \mathbb{E}_{(x,y) \sim \mathcal{D}} [\phi^\lambda(h(Mx), y)],$$

we have the following.<sup>3</sup>

**Lemma 6** Let  $\mathcal{M}$  be any class of weighting metrics on the feature space  $X = \mathbb{R}^D$ . For any  $\gamma > 0$ , let  $\mathcal{H}_{\sigma^\gamma}^{2\text{-net}}$  be a two layer feed-forward neural network base hypothesis class (as defined above) and  $\phi^\lambda$  be a classifier-based loss function that  $\lambda$ -Lipschitz in its first argument. Fix any sample size  $m$ , and let  $S_m$  be an i.i.d. sample of size  $m$  from an unknown bounded distribution  $\mathcal{D}$  on  $X \times \{0, 1\}$  (with bound  $B$ ). Then with probability at least  $1 - \delta$ ,

$$\sup_{M \in \mathcal{M}} [\text{err}_{\text{hypo}}^\lambda(M, \mathcal{D}) - \text{err}_{\text{hypo}}^\lambda(M, S_m)] \leq O\left(B\lambda\gamma\sqrt{\frac{d \ln(D/\delta)}{m}}\right),$$

<sup>3</sup>Since we know the functional form of the base hypothesis class  $\mathcal{H}$  (i.e., a two layer feed-forward neural net), we can provide a more precise bound than leaving it as  $\text{Fat}(\mathcal{H})$ .



where  $d$  is a uniform upperbound on the Frobenius norm of the quadratic form of weighting metrics in  $\mathcal{M}$ , that is,  $\sup_{M \in \mathcal{M}} \|M^\top M\|_F^2 \leq d$ .

### 4.3 Automatically Adapting to Intrinsic Complexity

Note that while Lemmas 5 and 6 provide a sample complexity bound that is tuned to the metric learning complexity of a given dataset, these results are *not* useful directly since one cannot select the correct norm bounded class  $\mathcal{M}$  a priori (as the underlying distribution  $\mathcal{D}$  is unknown).

Fortunately, by considering an appropriate sequence of norm-bounded classes of weighting metrics, we can provide a uniform bound that *automatically adapts* to the intrinsic complexity of the unknown underlying data distribution  $\mathcal{D}$ . In particular, we have the following.

**Corollary 7** *Fix any  $m$ , and let  $S_m$  be an i.i.d. sample of size  $m$  from an unknown bounded distribution  $\mathcal{D}$  (with bound  $B$ ). Define  $\mathcal{M}^d := \{M \mid \|M^\top M\|_F^2 \leq d\}$ , and consider the nested sequence of weighting metric class  $\mathcal{M}^1 \subset \mathcal{M}^2 \subset \dots$ . Let  $\mu_d$  be any non-negative measure across the sequence  $\mathcal{M}^d$  such that  $\sum_d \mu_d = 1$  (for  $d = 1, 2, \dots$ ). Then for any  $\lambda \geq 0$ , with probability at least  $1 - \delta$ , for all  $d = 1, 2, \dots$ , and all  $M^d \in \mathcal{M}^d$ ,*

$$[\text{err}^\lambda(M^d, \mathcal{D}) - \text{err}^\lambda(M^d, S_m)] \leq O\left(C \cdot B\lambda \sqrt{\frac{d \ln(1/\delta \mu_d)}{m}}\right), \quad (3)$$

where  $C := B$  for distance-based error, or  $C := \gamma\sqrt{\ln D}$  for classifier-based error (with base hypothesis class  $\mathcal{H}_{\sigma^\gamma}^{2\text{-net}}$ ).

In particular, for a data distribution  $\mathcal{D}$  that has metric learning complexity at most  $d \in \mathbb{N}$ , if there are  $m \geq \Omega\left(\frac{d(CB\lambda)^2 \ln(1/\delta \mu_d)}{\epsilon^2}\right)$  samples, then with probability at least  $1 - \delta$

$$[\text{err}^\lambda(M_m^{\text{reg}}, \mathcal{D}) - \text{err}^\lambda(M^*, \mathcal{D})] \leq O(\epsilon),$$

for  $M_m^{\text{reg}} := \text{argmin}_{M \in \mathcal{M}} [\text{err}^\lambda(M, S_m) + \Lambda_M d_M]$ , where  $\Lambda_M := CB\lambda \sqrt{\ln\left(\frac{1}{\delta \mu_{d_M}}\right)/m}$  and  $d_M := \lceil \|M^\top M\|_F^2 \rceil$ .

Observe that the measure  $(\mu_d)$  above encodes our prior belief on the complexity class  $\mathcal{M}^d$  from which a target metric is selected by a metric learning algorithm given the training sample  $S_m$ . In absence of any prior beliefs,  $\mu_d$  can be simply set to  $1/D$  (for  $d = 1, \dots, D$ ) for unit spectral-norm weighting metrics.

Thus, for an unknown underlying data distribution  $\mathcal{D}$  with metric learning complexity  $d$ , with number of samples just proportional to  $d$ , we can find a good weighting metric.

This result also highlights that the generalization error of *any* weighting metric returned by an algorithm is proportional to the (smallest) norm-bounded class to which it belongs (cf. Eq. 3). If two metrics  $M_1$  and  $M_2$  have similar empirical errors on a given sample, but have different intrinsic complexities, then the expected risk of the two metrics can be considerably different. We expect the metric with lower intrinsic complexity to yield better generalization error. This partly explains the observed empirical success of various types of norm-regularized optimization criteria for finding the optimal weighting metric (Lim et al., 2013; Law et al., 2014).

Using this as a guiding principle, we can design an improved optimization criteria for metric learning problems that jointly minimizes the sample error and a Frobenius norm regularization penalty. In particular,

$$\min_{M \in \mathcal{M}} \text{err}(M, S_m) + \Lambda \|M^\top M\|_F^2 \quad (4)$$

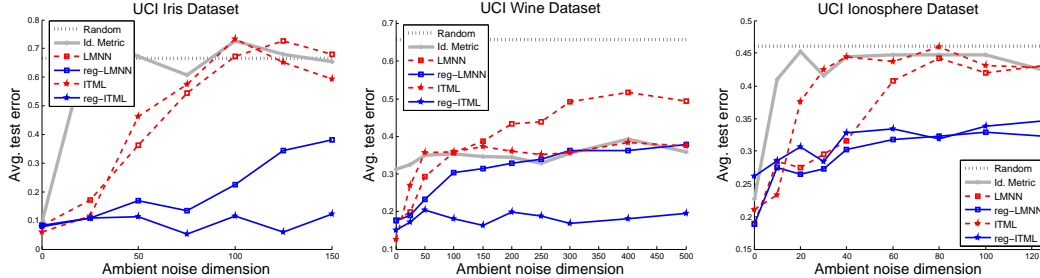


Figure 1: Nearest-neighbor classification performance of LMNN and ITML metric learning algorithms without regularization (dashed red lines) and with regularization (solid blue lines) on benchmark UCI datasets. The horizontal dotted line is the classification error of random label assignment drawn according to the class proportions, and solid gray line shows classification error of  $k$ -NN performance with respect to identity metric (no metric learning) for baseline reference.

for any error criteria ‘err’ used in a downstream prediction task of interest and a regularization hyper-parameter  $\Lambda$  proportional to  $m^{-1/2}$ . We explore the practical efficacy of this augmented optimization on some representative applications below.

## 5 Empirical Evaluation

Our analysis shows that the generalization error of metric learning can scale with the representation dimension, and regularization can help mitigate this by adapting to the intrinsic *metric learning complexity* of the given dataset. We want to explore to what degree these effects manifest in practice.

We select two popular metric learning algorithms, LMNN by Weinberger & Saul (2009) and ITML by Davis et al. (2007), that are designed to find metrics that improve nearest-neighbor classification quality. These algorithms have varying degrees of regularization built into their optimization criteria: LMNN implicitly regularizes the metric via its “large margin” criterion, while ITML allows for explicit regularization by letting the practitioners specify a “prior” weighting metric. We modified the LMNN optimization criteria as per Eq. (4) to also allow for an explicit norm-regularization controlled by the trade-off parameter  $\Lambda$ .

We can evaluate how the unregularized criteria (*i.e.*, unmodified LMNN, or ITML with the prior set to the identity matrix) compares to the regularized criteria (*i.e.*, modified LMNN with best  $\Lambda$ , or ITML with the prior set to a low-rank matrix).

**Datasets.** We use the UCI benchmark datasets for our experiments: IRIS (4 dim., 150 samples), WINE (13 dim., 178 samples) and IONOSPHERE (34 dim., 351 samples) datasets (Bache & Lichman, 2013). Each dataset has a fixed (unknown) intrinsic dimension; we can vary the representation dimension by augmenting each dataset with synthetic correlated noise of varying dimensions, simulating regimes where datasets contain large numbers of uninformative features.

Each UCI dataset is augmented with synthetic  $D$ -dimensional correlated noise as follows. We first sample a covariance matrix  $\Sigma_D$  from unit-scale Wishart distribution (that is, let  $A$  be a  $D \times D$  Gaussian random matrix with entry  $A_{ij} \sim N(0, 1)$  drawn i.i.d., and set  $\Sigma_D := A^T A$ ). Then each sample  $x_i$  from the dataset is appended independently by drawing noise vector  $x_\sigma \sim N(0, \Sigma_D)$ .

**Experimental setup.** We varied the ambient noise dimension  $D$  between 0 and 500 dimensions and

added it to the UCI datasets, creating the noise-augmented datasets. Each noise-augmented dataset was randomly split between 70% training, 10% validation, and 20% test samples.

We used the default settings for each algorithm. For regularized LMNN, we picked the best performing trade-off parameter  $\Lambda$  from  $\{0, 0.1, 0.2, \dots, 1\}$  on the validation set. For regularized ITML, we seeded with the rank-one discriminating metric, *i.e.*, we set the prior as the matrix with all zeros, except the diagonal entry corresponding to the most discriminating coordinate set to one.

All the reported results were averaged over 20 runs.

**Results.** Figure 1 shows the nearest-neighbor performance (with  $k = 3$ ) of LMNN and ITML on noise-augmented UCI datasets. Notice that the unregularized versions of both algorithms (dashed red lines) scale poorly when noisy features are introduced. As the number of uninformative features grows, the performance of both algorithms quickly degrades to that of classification performance in the original unweighted space with no metric learning (solid gray line), showing poor adaptability to the signal in the data.

Interestingly, neither of the unregularized algorithms performs consistently better than the other on datasets with high noise: ITML yields better results on WINE, whereas LMNN seems better for IONOSPHERE, and both algorithms yield similar performance on IRIS.

The regularized versions of both algorithms (solid blue lines) significantly improve the classification performance. Remarkably, regularized ITML shows almost no degradation in classification performance, even in very high noise regimes, demonstrating a strong robustness to noise.

These results underscore the value of regularization in metric learning, showing that regularization encourages adaptability to the intrinsic complexity and improved robustness to noise.

## 6 Discussion and Conclusion

Previous theoretical work on metric learning has focused almost exclusively on analyzing the generalization error of variants of the optimization criteria for the distance-based metric learning framework.

Jin et al. (2009), for instance, analyzed the generalization ability of regularized, convex-loss optimization criteria for pairwise distances via an algorithmic *stability* analysis. They derive an interesting sample complexity result that is sublinear in  $\sqrt{D}$  for datasets of representation dimension  $D$ . They discuss that the sample complexity can potentially be independent of  $D$ , but do not characterize specific instances or classes of problems where this may be possible.

Likewise, recent work by Bellet & Habrard (2012) uses algorithmic *robustness* to analyze the generalization ability for pairwise- and triplet-based distance metric learning. Their analysis relies on the existence of a partition of the input space, such that in each cell of the partition, the training loss and test loss does not deviate much (robustness criteria). Note that their sample complexity bound scales with the partition size, which in general can be exponential in the representation dimension.

Perhaps the works most similar to our approach are the sample complexity analyses by Bian & Tao (2011) and Cao et al. (2013). Bian & Tao (2011) analyze the consistency of the ERM criterion for metric learning. They show a  $O(m^{-1/2})$  rate of convergence for the ERM with  $m$  samples to the expected risk for thresholds on bounded convex losses for distance-based metric learning. Our upper-bound in Lemma 1 generalizes this result by considering arbitrary (possibly non-convex) distance-based Lipschitz losses and explicitly shows the dependence on the representation dimension  $D$ . Cao et al. (2013) provide an alternate analysis based on norm regularization of the weighting metric for distance-based metric learning. Their result parallels our norm-regularized criterion in

Lemma 5. While they focus on analyzing a specific optimization criterion – thresholds on the hinge loss with norm-regularization, our result holds for general Lipschitz losses.

It is worth emphasizing that none of these related works discuss the importance of or leverage the intrinsic structure in data for the metric learning problem. Our results in Section 4 formalize an intuitive notion of dataset’s intrinsic complexity for metric learning and show sample complexity rates that are finely tuned to this *metric learning complexity*.

The classifier-based framework we discuss has parallels with the kernel learning literature. The typical focus in kernel learning is to analyze the generalization ability of the hypothesis class of linear separators in general Hilbert spaces (Ying & Campbell, 2009; Cortes et al., 2010). Our work provides a complementary analysis for learning explicit linear transformations of the given representation space for arbitrary hypotheses classes.

Our theoretical analysis partly justifies the empirical success of norm-based regularization as well. Our empirical results show that such regularization not only helps in designing new metric learning algorithms (Lim et al., 2013; Law et al., 2014), but can even benefit existing metric learning algorithms in high-noise regimes.

## References

- Anthony, M. and Bartlett, P. *Neural network learning: Theoretical foundations*. Cambridge University Press, 1999.
- Bache, K. and Lichman, M. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Bartlett, P. and Mendelson, S. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research (JMLR)*, 3:463–482, 2002.
- Bellet, A. and Habrard, A. Robustness and generalization for metric learning. *CoRR*, abs/1209.1086, 2012. URL <http://arxiv.org/abs/1209.1086>.
- Bellet, A., Habrard, A., and Sebban, M. A survey on metric learning for feature vectors and structured data. *CoRR*, abs/1306.6709, 2014. URL <http://arxiv.org/abs/1306.6709>.
- Bian, W. and Tao, D. Learning a distance metric by empirical loss minimization. *International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1186–1191, 2011.
- Cao, Q., Guo, Z., and Ying, Y. Generalization bounds for metric and similarity learning. *CoRR*, abs/1207.5437, 2013. URL <http://arxiv.org/abs/1207.5437>.
- Cortes, C., Mohri, M., and Rostamizadeh, A. New generalization bounds for learning kernels. *International Conference on Machine Learning (ICML)*, 2010.
- Davis, J.V., Kulis, B., Jain, P., Sra, S., and Dhillon, I.S. Information-theoretic metric learning. *International Conference on Machine Learning (ICML)*, pp. 209–216, 2007.
- Guo, Z. and Ying, Y. Generalization classification via regularized similarity learning. *Neural Computation*, 26(3):497–552, 2014.
- Hornik, K., Stinchcombe, M., and White, H. Multilayer feedforward networks are universal approximators. *Neural Networks*, 4:359–366, 1989.

- Jin, R., Wang, S., and Zhou, Y. Regularized distance metric learning: Theory and algorithm. *Neural Information Processing Systems (NIPS)*, pp. 862–870, 2009.
- Law, M.T., Thome, N., and Cord, M. Fantope regularization in metric learning. *Computer Vision and Pattern Recognition (CVPR)*, 2014.
- Lim, D.K.H., McFee, B., and Lanckriet, G.R.G. Robust structural metric learning. *International Conference on Machine Learning (ICML)*, 2013.
- McFee, B. and Lanckriet, G.R.G. Metric learning to rank. *International Conference on Machine Learning (ICML)*, 2010.
- Shaw, B., Huang, B., and Jebara, T. Learning a distance metric from a network. *Neural Information Processing Systems (NIPS)*, 2011.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing, Theory and Applications*. 2010.
- Weinberger, K.Q. and Saul, L.K. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research (JMLR)*, 10:207–244, 2009.
- Xing, E.P., Ng, A.Y., Jordan, M.I., and Russell, S.J. Distance metric learning with application to clustering with side-information. *Neural Information Processing Systems (NIPS)*, pp. 505–512, 2002.
- Ying, Y. and Campbell, C. Generalization bounds for learning the kernel. *Conference on Computational Learning Theory (COLT)*, 2009.

## A Appendix: Various Proofs

### A.1 Proof of Lemma 1

Let  $\mathcal{P}$  be the probability measure induced by the random variable  $(\mathbf{X}, Y)$ , where  $\mathbf{X} := (x, x')$ ,  $Y := \mathbf{1}[y = y']$ , st.  $((x, y), (x', y')) \sim (\mathcal{D} \times \mathcal{D})$ .

Define function class

$$\mathcal{F} := \left\{ f_M : \mathbf{X} \mapsto \|M(x - x')\|^2 \mid \begin{array}{l} M \in \mathcal{M} \\ \mathbf{X} = (x, x') \in (X \times X) \end{array} \right\},$$

and consider any loss function  $\phi^\lambda(\rho, Y)$  that is  $\lambda$ -Lipschitz in the first argument. Then, we are interested in bounding the quantity

$$\sup_{f_M \in \mathcal{F}} \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{P}} [\phi^\lambda(f_M(\mathbf{X}), Y)] - \frac{1}{m} \sum_{i=1}^m \phi^\lambda(f_M(\mathbf{X}_i), Y_i),$$

where  $\mathbf{X}_i := (x_{1,i}, x_{2,i})$ ,  $Y_i := \mathbf{1}[y_{1,i} = y_{2,i}]$  from the paired sample  $S_m = \{((x_{1,i}, y_{1,i}), (x_{2,i}, y_{2,i}))\}_{i=1}^m$ .

Define  $\bar{x}_i := x_{1,i} - x_{2,i}$  for each  $\mathbf{X}_i = (x_{1,i}, x_{2,i})$ . Then, the Rademacher complexity<sup>4</sup> of our function class  $\mathcal{F}$  (with respect to the distribution  $\mathcal{P}$ ) is bounded, since (let  $\sigma_1, \dots, \sigma_m$  denote independent uniform  $\{\pm 1\}$ -valued random variables)

$$\begin{aligned} \mathcal{R}_m(\mathcal{F}, \mathcal{P}) &:= \mathbb{E}_{\mathbf{X}_i, \sigma_i} \left[ \sup_{f_M \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \sigma_i f_M(\mathbf{X}_i) \right] \\ &= \frac{1}{m} \mathbb{E}_{\mathbf{X}_i, \sigma_i} \sup_{i \in [m]} \sup_{M \in \mathcal{M}} \left[ \sum_{i=1}^m \sigma_i \bar{x}_i^\top M^\top M \bar{x}_i \right] \\ &= \frac{1}{m} \mathbb{E}_{\mathbf{X}_i, \sigma_i} \sup_{i \in [m]} \sup_{\substack{M \in \mathcal{M}, \text{ s.t.} \\ [a^{jk}]_{jk} := M^\top M}} \left[ \sum_{j,k} a^{jk} \sum_{i=1}^m \sigma_i \bar{x}_i^j \bar{x}_i^k \right] \\ &\leq \frac{1}{m} \mathbb{E}_{\mathbf{X}_i, \sigma_i} \sup_{i \in [m]} \sup_{M \in \mathcal{M}} \left[ \|M^\top M\|_F \left( \sum_{j,k} \left( \sum_{i=1}^m \sigma_i \bar{x}_i^j \bar{x}_i^k \right)^2 \right)^{1/2} \right] \\ &\leq \frac{\sqrt{D}}{m} \mathbb{E}_{\mathbf{X}_i, i \in [m]} \left( \mathbb{E}_{\sigma_i, i \in [m]} \sum_{j,k} \left( \sum_{i=1}^m \sigma_i \bar{x}_i^j \bar{x}_i^k \right)^2 \right)^{1/2} \\ &= \frac{\sqrt{D}}{m} \mathbb{E}_{\mathbf{X}_i, i \in [m]} \left( \sum_{j,k} \sum_{i=1}^m \left( \bar{x}_i^j \right)^2 \left( \bar{x}_i^k \right)^2 \right)^{1/2} \\ &= \frac{\sqrt{D}}{m} \mathbb{E}_{\mathbf{X}_i, i \in [m]} \left( \sum_{i=1}^m \|\bar{x}_i\|^4 \right)^{1/2} \\ &= \frac{\sqrt{D}}{m} \mathbb{E}_{(x_i, x'_i) \sim (\mathcal{D} \times \mathcal{D} \mid x \times \mathcal{D} \mid x), i \in [m]} \left( \sum_{i=1}^m \|x_i - x'_i\|^4 \right)^{1/2} \end{aligned}$$

<sup>4</sup>See the definition of Rademacher complexity in the statement of Lemma 8.

$$\begin{aligned}
&\leq \sqrt{\frac{D}{m}} \left( \mathbb{E}_{(x,x') \sim (\mathcal{D}|_X \times \mathcal{D}|_X)} \|x - x'\|^4 \right)^{1/2} \\
&\leq 4B^2 \sqrt{\frac{D}{m}},
\end{aligned}$$

where the second inequality is by noting that  $\sup_{M \in \mathcal{M}} \|M^\top M\|_F \leq \sqrt{D}$  for the class of weighting metrics  $\mathcal{M} := \{M \mid M \in \mathbb{R}^{D \times D}, \sigma_{\max}(M) = 1\}$ .

Recall that  $\mathcal{D}$  has bounded support (with bound  $B$ ). Thus, by noting that  $\phi^\lambda$  is  $8B^2$  bounded function that is  $\lambda$ -Lipschitz in the first argument, we can apply Lemma 8 and get the desired uniform deviation bound. ■

**Lemma 8 [Rademacher complexity of bounded Lipschitz loss functions Bartlett & Mendelson (2002)]** *Let  $\mathcal{D}$  be a fixed unknown distribution over  $X \times \{-1, 1\}$ , and let  $S_m$  be an i.i.d. sample of size  $m$  from  $\mathcal{D}$ . Given a hypothesis class  $\mathcal{H} \subset \mathbb{R}^X$  and a loss function  $\ell : \mathbb{R} \times \{-1, 1\} \rightarrow \mathbb{R}$ , such that  $\ell$  is  $c$ -bounded, and is  $\lambda$ -Lipschitz in the first argument, that is,  $\sup_{(y', y) \in \mathbb{R} \times \{-1, 1\}} |\ell(y', y)| \leq c$ , and  $|\ell(y', y) - \ell(y'', y)| \leq \lambda|y' - y''|$ , we have the following:*

*for any  $0 < \delta < 1$ , with probability at least  $1 - \delta$ , every  $h \in \mathcal{H}$  satisfies*

$$\text{err}(\ell \circ h, \mathcal{D}) \leq \text{err}(\ell \circ h, S_m) + 2\lambda \mathcal{R}_m(\mathcal{H}, \mathcal{D}) + c \sqrt{\frac{2 \ln(1/\delta)}{m}},$$

where

- $\text{err}(\ell \circ h, \mathcal{D}) := \mathbb{E}_{x, y \sim \mathcal{D}} [\ell(h(x), y)]$ ,
- $\text{err}(h, S_m) := \frac{1}{m} \sum_{(x_i, y_i) \in S_m} \ell(h(x_i), y_i)$ ,
- $\mathcal{R}_m(\mathcal{H}, \mathcal{D})$  is the Rademacher complexity of the function class  $\mathcal{H}$  with respect to the distribution  $\mathcal{D}$  given  $m$  i.i.d. samples, and is defined as:

$$\mathcal{R}_m(\mathcal{H}, \mathcal{D}) := \mathbb{E}_{\substack{x_i \sim \mathcal{D}|_X, \\ \sigma_i \sim \text{unif}\{\pm 1\}, \\ i \in [m]}} \left[ \sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \sigma_i h(x_i) \right],$$

where  $\sigma_i$  are independent uniform  $\{\pm 1\}$ -valued random variables.

## A.2 Proof of Lemma 2

We shall exhibit a finite class of bounded support distributions  $\mathfrak{D}$ , such that if  $\mathcal{D}$  is chosen uniformly at random from  $\mathfrak{D}$ , the expectation (over the random choice of  $\mathcal{D}$ ) of the probability of failure (that is, generalization error of the metric returned by  $\mathcal{A}$  compared to that of the optimal metric exceeds the specified tolerance level  $\epsilon$ ) is at least  $\delta$ . This implies that for some distribution in  $\mathfrak{D}$  the probability of failure is at least  $\delta$  as well.

Let  $\Delta_D := \{x_0, \dots, x_D\}$  be a set of  $D + 1$  points that from the vertices of a regular unit-simplex from the underlying space  $X = \mathbb{R}^D$  as per Definition 1 (see below). For a fixed parameter  $0 < \alpha < 1$  (exact value determined later), define  $\mathfrak{D}$  as the class of all distributions  $\mathcal{D}$  on  $X \times \{0, 1\}$  such that:

- $\mathcal{D}$  assigns zero probability to all sets not intersecting  $\Delta_D \times \{0, 1\}$ .

- for each  $i = 0, \dots, D$ , either
  - $\mathbf{P}[(x_i, 1)] = (1 + \sqrt{\alpha})/2$  and  $\mathbf{P}[(x_i, 0)] = (1 - \sqrt{\alpha})/2$ , or
  - $\mathbf{P}[(x_i, 1)] = (1 - \sqrt{\alpha})/2$  and  $\mathbf{P}[(x_i, 0)] = (1 + \sqrt{\alpha})/2$ .

For concreteness, we shall use a specific instantiation of  $\phi_{L,U}^\lambda$  in  $\text{err}_{\text{dist}}^\lambda$  with  $U = 0$ ,  $L = 4/D$  and  $\lambda = D/4$ .

**Proof overview.** We first show, by the construction of the distributions under consideration in  $\mathfrak{D}$ , the sample error and the generalization error minimizing metrics over any  $\mathcal{D} \in \mathfrak{D}$  belong to a restricted class of weighting matrices (Eq. 5). We then make a second simplification by noting that finding these (sample- and generalization-) error minimizing metrics (in the restricted class) is equivalent to solving a binary classification problem (Eq. 6). This reduction to binary classification enables us to use VC-style lower bounding techniques to give a lower bound on the sample complexity. We now fill in the details.

Consider a subset of weighting metrics  $\mathcal{M}_{0.1}$  that map points in  $\Delta_D$  to exactly one of two possible points that are (squared) distance at least  $4/D$  apart, that is,

$$\begin{aligned} \mathcal{M}_{0.1} := \{M \mid M \in \mathcal{M}, \exists z_0, z_1 \in \mathbb{R}^D, \forall x \in \Delta_D, \\ Mx \in \{z_0, z_1\} \text{ and } \|z_0 - z_1\|^2 \geq 4/D\}. \end{aligned}$$

Now pick any  $\mathcal{D} \in \mathfrak{D}$ , let  $S_m$  be an i.i.d. paired sample from  $\mathcal{D}$ . Observe that both the sample-based and the distribution-based error minimizing weighting metric from  $\mathcal{M}$  on  $\mathcal{D}$  also belongs to  $\mathcal{M}_{0.1}$ . That is, (c.f. Lemma 10)

$$\begin{aligned} \text{argmin}_{M \in \mathcal{M}} \text{err}_{\text{dist}}(M, \mathcal{D}) &= \text{argmin}_{M \in \mathcal{M}_{0.1}} \text{err}_{\text{dist}}(M, \mathcal{D}) \\ \text{argmin}_{M \in \mathcal{M}} \text{err}_{\text{dist}}(M, S_m) &= \text{argmin}_{M \in \mathcal{M}_{0.1}} \text{err}_{\text{dist}}(M, S_m). \end{aligned} \quad (5)$$

**A reduction to binary classification on product space.** For each  $M \in \mathcal{M}_{0.1}$ , we associate a classifier  $f_M : (\Delta_D \times \Delta_D) \rightarrow \{0, 1\}$  defined as  $(x_i, x_j) \mapsto \mathbf{1}[Mx_i = Mx_j]$ . Now, consider the probability measure  $\mathcal{P}$  induced by the random variable  $(\mathbf{X}, Y)$ , where  $\mathbf{X} := (x, x')$ ,  $Y := \mathbf{1}[y = y']$ , s.t.  $((x, y), (x', y')) \sim (\mathcal{D}|_{(\Delta_D \times \{0,1\})} \times \mathcal{D}|_{(\Delta_D \times \{0,1\})})$ . It is easy to check that for all  $M \in \mathcal{M}_{0.1}$

$$\begin{aligned} \text{err}_{\text{dist}}^\lambda(M, \mathcal{D}) &= \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{P}} [\mathbf{1}[f_M(\mathbf{X}) \neq Y]] \\ \text{err}_{\text{dist}}^\lambda(M, S_m) &= \frac{1}{m} \sum_{((x,y),(x',y')) \in S_m} \mathbf{1}[f_M((x, x')) \neq \mathbf{1}[y = y']]. \end{aligned} \quad (6)$$

Define

$$\begin{aligned} \eta(\mathbf{X}) &:= \mathbf{P}_{Y \sim \mathcal{P}|_{Y|\mathbf{X}}} [Y = 1 | \mathbf{X}] \\ &= \mathbf{P}_{(y,y') \sim (\mathcal{D} \times \mathcal{D})|_{(y,y')|(x,x')}} [y = y' | x, x'] \\ &= \begin{cases} \frac{1}{2} + \frac{\alpha}{2} & \text{if } \mathbf{P}(y|x) = \mathbf{P}(y'|x') \\ \frac{1}{2} - \frac{\alpha}{2} & \text{if } \mathbf{P}(y|x) \neq \mathbf{P}(y'|x') \end{cases}. \end{aligned} \quad (7)$$

Observe that  $\eta(\mathbf{X})$  is the Bayes error rate at  $\mathbf{X}$  for distribution  $\mathcal{P}$ . Since, by construction of  $\mathcal{M}_{0.1}$ , the class  $\{f_M\}_{M \in \mathcal{M}_{0.1}}$  contains a classifier that achieves the Bayes error rate, the optimal classifier



$f^* := \operatorname{argmin}_{f_M} \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{P}} \mathbf{1}[f_M(\mathbf{X}) \neq Y]$  necessarily has  $f^*(\mathbf{X}) = \mathbf{1}[\eta(\mathbf{X}) > \frac{1}{2}]$  (for all  $\mathbf{X}$ ). Then, for any  $f_M$ ,

$$\begin{aligned}
& \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{P}} [\mathbf{1}[f_M(\mathbf{X}) \neq Y]] - \mathbb{E}_{(\mathbf{X}, Y) \sim \mathcal{P}} [\mathbf{1}[f^*(\mathbf{X}) \neq Y]] \\
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{P}|\mathbf{x}} [\eta(\mathbf{X})(\mathbf{1}[f^*(\mathbf{X}) = 1] - \mathbf{1}[f_M(\mathbf{X}) = 1]) \\
&\quad + (1 - \eta(\mathbf{X}))(\mathbf{1}[f^*(\mathbf{X}) = 0] - \mathbf{1}[f_M(\mathbf{X}) = 0])] \\
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{P}|\mathbf{x}} [(2\eta(\mathbf{X}) - 1)(\mathbf{1}[f^*(\mathbf{X}) = 1] - \mathbf{1}[f_M(\mathbf{X}) = 1])] \\
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{P}|\mathbf{x}} [2|\eta(\mathbf{X}) - 1/2| \cdot \mathbf{1}[f_M(\mathbf{X}) \neq f^*(\mathbf{X})]] \\
&= \frac{2\alpha}{(D+1)^2} \sum_{i>j} [\mathbf{1}[f_M((x_i, x_j)) \neq f^*((x_i, x_j))]], \tag{8}
\end{aligned}$$

where (i) the second to last equality is by noting that  $f^*(\mathbf{X}) \neq 1 \iff \eta(\mathbf{X}) \leq 1/2$ , and (ii) the last equality is by noting Eq. (7),  $f_M((x_i, x_i)) = f^*((x_i, x_i)) = 1$  for all  $i$  and  $f((x_i, x_j)) = f((x_j, x_i))$  for all  $f$ . For notational simplicity, we shall define  $\mathbf{X}_{i,j} := (x_i, x_j)$ .

Now, for a given paired sample  $S_m$ , let  $N(S_m) := (N_i)_i$  (for all  $0 \leq i \leq D$ ), where  $N_i$  is the number of occurrences of the point  $x_i$  in  $S_m$ . Then for any  $f_M$ ,

$$\begin{aligned}
& \mathbb{E}_{S_m} \left[ \frac{1}{(D+1)^2} \sum_{i>j} \mathbf{1}[f_M(\mathbf{X}_{i,j}) \neq f^*(\mathbf{X}_{i,j})] \right] \\
&= \frac{1}{(D+1)^2} \sum_{i>j} \mathbf{P}_{S_m}[f_M(\mathbf{X}_{i,j}) \neq f^*(\mathbf{X}_{i,j})] \\
&= \frac{1}{(D+1)^2} \sum_{i>j} \sum_{N \in \mathbb{N}^{D+1}} \mathbf{P}_{S_m}[f_M(\mathbf{X}_{i,j}) \neq f^*(\mathbf{X}_{i,j}) | N(S_m) = N] \cdot \mathbf{P}[N(S_m) = N] \\
&= \frac{1}{(D+1)^2} \sum_{N \in \mathbb{N}^{D+1}} \mathbf{P}[N(S_m) = N] \cdot \sum_{i>j} \mathbf{P}_{S_m}[f_M(\mathbf{X}_{i,j}) \neq f^*(\mathbf{X}_{i,j}) | N_i, N_j] \\
&\geq \frac{1}{(D+1)^2} \sum_{N \in \mathbb{N}^{D+1}} \mathbf{P}[N(S_m) = N] \cdot \sum_{i>j} \frac{1}{4} \left( 1 - \sqrt{1 - \exp \left( \frac{-(\max\{N_i, N_j\} + 1)\alpha^2}{1 - \alpha^2} \right)} \right) \\
&\geq \frac{1}{4} \frac{D}{D+1} \left( 1 - \sqrt{1 - \exp \left( \frac{-((2m/(D+1)) + 1)\alpha^2}{1 - \alpha^2} \right)} \right) \\
&\geq \frac{1}{8} \left( 1 - \sqrt{1 - \exp \left( \frac{-((2m/(D+1)) + 1)\alpha^2}{1 - \alpha^2} \right)} \right),
\end{aligned}$$

where (i) the first inequality is by applying Lemma 11, (ii) the second inequality is by assuming WLOG  $N_i \geq N_j$ , and noting that the expression above is convex in  $N_i$  so one can apply Jensen's inequality and by observing that  $\mathbb{E}[N_i] = 2m/(D+1)$  and that there are total  $D(D+1)$  summands for  $i > j$ , and (iii) the last inequality is by noting that  $D \geq 1$ . Now, let  $B$  denote the r.h.s. quantity above. Then by recalling that for any  $[0, 1]$ -valued random variable  $Z$ ,  $\mathbf{P}(Z > \gamma) > \mathbb{E}Z - \gamma$  (for all  $0 < \gamma < 1$ ), we have

$$\mathbf{P}_{S_m} \left[ \frac{1}{(D+1)^2} \sum_{i>j} \mathbf{1}[f_M((x_i, x_j)) \neq f^*((x_i, x_j))] > \gamma B \right] > (1 - \gamma)B.$$

Or equivalently, by combining Eqs. (5), (6) and (8), we have

$$\mathbb{E}_{\mathcal{D} \sim \text{unif}(\mathfrak{D})} \mathbf{P}_{S_m \sim \mathcal{D}} \left[ \text{err}_{\text{dist}}(\mathcal{A}(S_m), \mathcal{D}) - \text{err}_{\text{dist}}(M_{\mathcal{D}}^*, \mathcal{D}) > 2\alpha\gamma B \right] > (1 - \gamma)B,$$

where  $M_{\mathcal{D}}^* := \text{argmin}_{M \in \mathcal{M}} \text{err}_{\text{dist}}(M, \mathcal{D})$  and  $\mathcal{A}(S_m)$  is any metric returned by empirical error minimizing algorithm. Now, if (cond. 1)  $B \geq \delta/1 - \gamma$  and (cond. 2)  $\epsilon \leq 2\gamma\alpha B$ , it follows that for some  $\mathcal{D} \in \mathfrak{D}$

$$\mathbf{P}_{S_m \sim \mathcal{D}} \left[ \text{err}_{\text{dist}}(\mathcal{A}(S_m), \mathcal{D}) - \text{err}_{\text{dist}}(M_{\mathcal{D}}^*, \mathcal{D}) > \epsilon \right] > \delta. \quad (9)$$

Now, to satisfy cond. 1 & 2, we shall select  $\gamma = 1 - 16\delta$ . Then cond. 1 follows if

$$m \leq \frac{(D+1)}{2} \left( \frac{1 - \alpha^2}{\alpha^2} \ln(4/3) - 1 \right).$$

Choosing parameter  $\alpha = 8\epsilon/\gamma$  (and by noting  $B \geq 1/16$  by cond. 1 for choice of  $\gamma$  and  $m$ ), cond. 2 is satisfied as well. Hence,

$$m \leq \frac{(D+1)}{2} \left( \frac{(1 - 16\delta)^2 - (8\epsilon)^2}{64\epsilon^2} \ln(4/3) - 1 \right)$$

implies Eq. (9). Moreover, if  $0 < \epsilon, \delta < 1/64$  then  $m \leq \frac{(D+1)}{512\epsilon^2}$  would suffice. ■

**Definition 1** Define  $n + 1$  vectors  $\Delta_n = \{v_0, \dots, v_n\}$ , with each  $v_i \in \mathbb{R}^n$  as

$$\begin{aligned} v_{0,j} &= \frac{-1}{\sqrt{n}} && \text{for } 1 \leq j \leq n \\ v_{i,j} &= \begin{cases} \frac{(n-1)\sqrt{n+1}+1}{n\sqrt{n}} & \text{if } i = j \\ \frac{-(\sqrt{n+1}-1)}{n\sqrt{n}} & \text{otherwise} \end{cases} && \text{for } 1 \leq i, j \leq n \end{aligned}$$

**Fact 9 [properties of vertices of a regular  $n$ -simplex]** Let  $\Delta_n = \{v_0, \dots, v_n\}$  be a set of  $n + 1$  vectors in  $\mathbb{R}^n$  as per Definition 1. Then,  $\Delta_n$  defines vertices of a regular  $n$ -simplex circumscribed in a unit  $(n - 1)$ -sphere, with

- (i)  $\|v_i\|^2 = 1$  (for all  $i$ ), and
- (ii)  $\|v_i - v_j\|^2 = 2(n + 1)/n$  (for  $i \neq j$ ).

Moreover, for any non-empty bi-partition of  $\Delta_n$  into  $\Delta_n^{(1)}$  and  $\Delta_n^{(2)}$  with  $|\Delta_n^{(1)}| = k$  and  $|\Delta_n^{(2)}| = n + 1 - k$ , define  $a^{(1)}$  and  $a^{(2)}$  the means (centroids) of the points in  $\Delta_n^{(1)}$  and  $\Delta_n^{(2)}$  respectively. Then, we also have

- (i)  $(a^{(1)} - a^{(2)}) \cdot (a^{(i)} - v_j) = 0$  (for  $i \in \{1, 2\}$ , and  $v_j \in \Delta_n^{(i)}$ ).
- (ii)  $\|a^{(1)} - a^{(2)}\|^2 = \frac{(n+1)^2}{kn(n+1-k)} \geq \frac{4}{n}$ , for  $1 \leq k \leq n$ .

**Lemma 10** Let  $\Delta_D$  be a set of  $D + 1$  points  $\{X_0, \dots, X_D\}$  in  $\mathbb{R}^D$  as per Definition 1, and let  $\mathcal{D}$  be an arbitrary distribution over  $\Delta_D \times \{0, 1\}$ . Define  $P_i := \mathbf{1}[\mathbf{P}_{\mathcal{D}}[(X_i, 1)] > 1/2]$ . Define  $\Pi := \{\pi : \Delta_D \rightarrow \mathbb{R}^D\}$  be the collection of all functions that maps points in  $\Delta_D$  to arbitrary points in  $\mathbb{R}^D$ . Define

$$f((x, y), (x', y'); \pi) := \begin{cases} \min\{1, \frac{D}{4} \|\pi(x) - \pi(x')\|^2\} & \text{if } y = y' \\ \min\{1, [1 - \frac{D}{4} \|\pi(x) - \pi(x')\|^2]_+\} & \text{if } y \neq y' \end{cases}.$$

Let  $\mathcal{E}(\pi) := \mathbb{E}_{(x, y), (x', y') \sim \mathcal{D} \times \mathcal{D}}[f((x, y), (x', y'); \pi)]$  and  $\mathcal{E}^* := \inf_{\pi} \mathcal{E}(\pi)$ . Then, for any  $\bar{\pi} \in \Pi$  such that

- (i)  $\bar{\pi}(X_i) = \bar{\pi}(X_j)$ , if  $P_i = P_j$
- (ii)  $\|\bar{\pi}(X_i) - \bar{\pi}(X_j)\|^2 \geq \frac{4}{D}$ , if  $P_i \neq P_j$ ,

we have that  $\mathcal{E}(\bar{\pi}) = \mathcal{E}^*$ . Moreover, define  $\bar{A}$  as

- $\bar{A} := \frac{A_1 - A_0}{\|A_1 - A_0\|}$ , where  $A_0 := \text{mean}(X_i)$  such that  $P_i = 0$ , and  $A_1 := \text{mean}(X_i)$  such that  $P_i = 1$  (if exists at least one  $P_i = 0$  and at least one  $P_i = 1$ ).
- $\bar{A} := 0$ , i.e. the zero vector in  $\mathbb{R}^D$  (otherwise).

And let  $M$  be a  $D \times D$  matrix (with  $\sigma_{\max}(M) = 1$ ) defined as

$$M := \bar{A}\bar{A}^\top.$$

Then the map  $\pi_M : x \mapsto Mx$  constitutes a map that satisfies conditions (i) and (ii) and thus  $\mathcal{E}(\pi_M) = \mathcal{E}^*$ .

*Proof.* The proof follows from the geometric properties of  $\Delta_D$  and Fact 9. ■

**Lemma 11** Given two random variables  $\alpha_1$  and  $\alpha_2$ , each uniformly distributed on  $\{\alpha_-, \alpha_+\}$  independently, where  $\alpha_- = 1/2 - \epsilon/2$  and  $\alpha_+ = 1/2 + \epsilon/2$  with  $0 < \epsilon < 1$ . Suppose that  $\xi_1^1, \dots, \xi_m^1$  and  $\xi_1^2, \dots, \xi_m^2$  are two i.i.d. sequences of  $\{0, 1\}$ -valued random variables with  $\mathbf{P}(\xi_i^1 = 1) = \alpha_1$  and  $\mathbf{P}(\xi_i^2 = 1) = \alpha_2$  for all  $i$ . Then, for any likelihood maximizing function  $f$  from  $\{0, 1\}^m$  to  $\{\alpha_-, \alpha_+\}$  that estimates the bias  $\alpha_1$  and  $\alpha_2$  from the samples,

$$\begin{aligned} & \mathbf{P}\left[(f(\xi_1^1, \dots, \xi_m^1) \neq \alpha_1 \text{ and } f(\xi_1^2, \dots, \xi_m^2) = \alpha_2), \right. \\ & \left. \text{or } (f(\xi_1^1, \dots, \xi_m^1) = \alpha_1 \text{ and } f(\xi_1^2, \dots, \xi_m^2) \neq \alpha_2)\right] > \frac{1}{4} \left(1 - \sqrt{1 - \exp\left(\frac{-2\lceil m/2 \rceil \epsilon^2}{1 - \epsilon^2}\right)}\right). \end{aligned}$$

*Proof.* Note that

$$\begin{aligned} & \mathbf{P}\left[(f(\xi_1^1, \dots, \xi_m^1) \neq \alpha_1 \text{ and } f(\xi_1^2, \dots, \xi_m^2) = \alpha_2), \text{ or } (f(\xi_1^1, \dots, \xi_m^1) = \alpha_1 \text{ and } f(\xi_1^2, \dots, \xi_m^2) \neq \alpha_2)\right] \\ &= \mathbf{P}[f(\xi_1^1, \dots, \xi_m^1) \neq \alpha_1] \cdot \mathbf{P}[f(\xi_1^2, \dots, \xi_m^2) = \alpha_2] + \mathbf{P}[f(\xi_1^1, \dots, \xi_m^1) = \alpha_1] \cdot \mathbf{P}[f(\xi_1^2, \dots, \xi_m^2) \neq \alpha_2] \\ &\geq \frac{1}{2} \mathbf{P}[f(\xi_1^1, \dots, \xi_m^1) \neq \alpha_1] + \frac{1}{2} \mathbf{P}[f(\xi_1^2, \dots, \xi_m^2) \neq \alpha_2] \\ &> \frac{1}{4} \left(1 - \sqrt{1 - \exp\left(\frac{-2\lceil m/2 \rceil \epsilon^2}{1 - \epsilon^2}\right)}\right), \end{aligned}$$

where the first inequality is by noting that a likelihood maximizing  $f$  will select the correct bias better than random (which has probability 1/2), and the second inequality is by applying Lemma 12. ■

**Lemma 12 [Lemma 5.1 of Anthony & Bartlett (1999)]** Suppose that  $\alpha$  is a random variable uniformly distributed on  $\{\alpha_-, \alpha_+\}$ , where  $\alpha_- = 1/2 - \epsilon/2$  and  $\alpha_+ = 1/2 + \epsilon/2$ , with  $0 < \epsilon < 1$ . Suppose that  $\xi_1, \dots, \xi_m$  are i.i.d.  $\{0, 1\}$ -valued random variables with  $\mathbf{P}(\xi_i = 1) = \alpha$  for all  $i$ . Let  $f$  be a function from  $\{0, 1\}^m$  to  $\{\alpha_-, \alpha_+\}$ . Then

$$\mathbf{P}[f(\xi_1, \dots, \xi_m) \neq \alpha] > \frac{1}{4} \left( 1 - \sqrt{1 - \exp\left(\frac{-2\lceil m/2 \rceil \epsilon^2}{1 - \epsilon^2}\right)} \right).$$

### A.3 Proof of Lemma 3

For any  $M \in \mathcal{M}$  define real-valued hypothesis class on domain  $X$  as  $\mathcal{H}_M := \{x \mapsto h(Mx) : h \in \mathcal{H}\}$  and define

$$\mathcal{F} := \{x \mapsto h(Mx) : M \in \mathcal{M}, h \in \mathcal{H}\} = \bigcup_M \mathcal{H}_M.$$

Observe that a uniform convergence of errors induced by the functions in  $\mathcal{F}$  implies convergence of the class of weighted matrices as well.

Now for any domain  $X$ , real-valued hypothesis class  $\mathcal{G} \subset [0, 1]^X$ , margin  $\gamma > 0$ , and a sample  $S \subset X$ , define

$$\text{cov}_\gamma(\mathcal{G}, S) := \left\{ C \subset \mathcal{G} \mid \begin{array}{l} \forall g \in \mathcal{G}, \exists g' \in C, \\ \max_{s \in S} |g(s) - g'(s)| \leq \gamma \end{array} \right\}$$

as the set of  $\gamma$ -covers of  $S$  by  $\mathcal{G}$ . Let  $\gamma$ -covering number of  $\mathcal{G}$  for any integer  $m > 0$  be defined as

$$\mathcal{N}_\infty(\gamma, \mathcal{G}, m) := \max_{S \subset X: |S|=m} \min_{C \in \text{cov}_\gamma(\mathcal{G}, S)} |C|,$$

with the minimizing cover  $C$  called as the minimizing  $(\gamma, m)$ -cover of  $\mathcal{G}$

Now, for the given  $\gamma$ , we will first estimate the  $\gamma$ -covering number of  $\mathcal{F}$ , that is,  $\mathcal{N}_\infty(\gamma, \mathcal{F}, m)$ .

For any  $M \in \mathcal{M}$ , let  $H_M$  be the minimizing  $(\gamma/2, m)$ -cover of  $\mathcal{H}_M$ . Note that  $|H_M| = \mathcal{N}_\infty(\gamma/2, \mathcal{H}_M, m) \leq \mathcal{N}_\infty(\gamma/2, \mathcal{H}, m)$  (because  $MX \subset X$ ).

Now let  $\mathcal{M}_\epsilon$  be an  $\epsilon$ -spectral cover of  $\mathcal{M}$  (that is, for every  $M \in \mathcal{M}$ , exists  $M' \in \mathcal{M}_\epsilon$  such that  $\sigma_{\max}(M - M') \leq \epsilon$ ), and define

$$\bar{F}_\epsilon := \{x \mapsto h(Mx) : M \in \mathcal{M}_\epsilon, h \in \mathcal{H}_M\}.$$

Note that  $|\bar{F}_\epsilon| \leq |\mathcal{M}_\epsilon| |\mathcal{H}| \leq \mathcal{N}_\infty(\gamma/2, \mathcal{H}, m) (1 + 2D/\epsilon)^{D^2}$  (c.f. Lemma 13). Observe that  $\bar{F}_\epsilon$  is a  $(\gamma/2 + B\lambda\epsilon)$ -cover of  $\mathcal{F}$ , since (i) for any  $f \in \mathcal{F}$  (formed by combining, say,  $M_0 \in \mathcal{M}$  and  $h_0 \in \mathcal{H}$ ), exists  $\bar{f} \in \bar{F}_\epsilon$ , namely the  $\bar{f}$  formed by  $\bar{M}_0$  such that  $\sigma_{\max}(M_0 - \bar{M}_0) \leq \epsilon$ , and (ii)  $\bar{h}_0 \in \mathcal{H}_{\bar{M}_0}$  such that  $|h_0(\bar{M}_0 x) - \bar{h}_0(\bar{M}_0 x)| \leq \gamma/2$  (for all  $x \in X$ ). So, (for any  $x \in X$ )

$$\begin{aligned} |f(x) - \bar{f}(x)| &= |h_0(M_0 x) - \bar{h}_0(\bar{M}_0 x)| \\ &\leq |h_0(M_0 x) - h_0(\bar{M}_0 x)| \\ &\quad + |h_0(\bar{M}_0 x) - \bar{h}_0(\bar{M}_0 x)| \\ &\leq \lambda \|M_0 x - \bar{M}_0 x\| + \gamma/2 \\ &\leq \lambda \sigma_{\max}(M_0 - \bar{M}_0) \|x\| + \gamma/2 \\ &\leq \lambda \epsilon B + \gamma/2. \end{aligned}$$

So, if we pick  $\epsilon = \min\{\frac{1}{2\lambda B}, \frac{\gamma}{2}\}$ , it follows that

$$\mathcal{N}_\infty(\gamma, \mathcal{F}, m) \leq |\bar{F}_\epsilon| \leq \mathcal{N}_\infty(\gamma/2, \mathcal{H}, m)(1 + 2D/\epsilon)^{D^2}.$$

By noting Lemmas 14 and 15, it follows that

$$\begin{aligned} \mathbf{P}_{S_m \sim \mathcal{D}} \left[ \exists f \in \mathcal{F} : \text{err}(f) \geq \text{err}_\gamma(f, S_m) + \alpha \right] \\ \leq 4 \left( 1 + \frac{2D}{\epsilon} \right)^{D^2} \left( \frac{128m}{\gamma^2} \right)^{\text{Fat}_{\gamma/16}(\mathcal{H}) \ln \left( \frac{32em}{\text{Fat}_{\gamma/16}(\mathcal{H})\gamma} \right)} e^{-\alpha^2 m/8}. \end{aligned}$$

The lemma follows by bounding this failure probability with at most  $\delta$ . ■

**Lemma 13 [ $\epsilon$ -spectral coverings of  $D \times D$  matrices]** *Let  $\mathcal{M} := \{M \mid M \in \mathbb{R}^{D \times D}, \sigma_{\max}(M) = 1\}$  be the set of matrices with unit spectral norm. Define  $\mathcal{M}_\epsilon$  as the  $\epsilon$ -cover of  $\mathcal{M}$ , that is, for every  $M \in \mathcal{M}$ , there exists  $M' \in \mathcal{M}_\epsilon$  such that  $\sigma_{\max}(M - M') \leq \epsilon$ . Then for all  $\epsilon > 0$ , there exists  $\mathcal{M}_\epsilon$  such that  $|\mathcal{M}_\epsilon| \leq \left(1 + \frac{2D}{\epsilon}\right)^{D^2}$ .*

*Proof.* Fix any  $\epsilon > 0$  and let  $\mathcal{N}_{\epsilon/D}$  be a minimal size  $(\epsilon/D)$ -cover of Euclidean unit ball  $\mathbf{B}_D$  in  $\mathbb{R}^D$ . That is, for any  $v \in \mathbf{B}_D$ , there exists  $v' \in \mathcal{N}_{\epsilon/D}$  such that  $\|v - v'\| \leq \epsilon/D$ . Using standard volume arguments (see e.g. proof of Lemma 5.2 of Vershynin (2010)), we know that  $|\mathcal{N}_{\epsilon/D}| \leq \left(1 + \frac{2D}{\epsilon}\right)^D$ . Define

$$\mathcal{M}_\epsilon := \left\{ M' \mid M' = [v'_1 \cdots v'_D] \in \mathbb{R}^{D \times D}, v'_i \in \mathcal{N}_{\epsilon/D} \right\}.$$

Then  $\mathcal{M}_\epsilon$  constitutes as an  $\epsilon$ -cover of  $\mathcal{M}$ , since for any  $M = [v_1 \cdots v_D] \in \mathcal{M}$  there exists  $M' = [v'_1 \cdots v'_D] \in \mathcal{M}_\epsilon$ , in particular  $M'$  such that  $\|v_i - v'_i\| \leq \epsilon/D$  (for all  $i$ ). Then

$$\sigma_{\max}(M - M') \leq \|M - M'\|_F = \sum_i \|v_i - v'_i\| \leq \epsilon.$$

Without loss of generality we can assume that each  $M' \in \mathcal{M}_\epsilon$ ,  $\sigma_{\max}(M') = 1$ . Moreover, by construction,  $|\mathcal{M}_\epsilon| \leq \left(1 + \frac{2D}{\epsilon}\right)^{D^2}$ . ■

**Lemma 14 [extension of Theorem 12.8 of Anthony & Bartlett (1999)]** *Let  $\mathcal{H}$  be a set of real functions from a domain  $X$  to the interval  $[0, 1]$ . Let  $\gamma > 0$ . Then for all  $m \geq 1$ ,*

$$\mathcal{N}_\infty(\gamma, \mathcal{H}, m) < c_0 (4m/\gamma^2)^{\text{Fat}_{\gamma/4}(\mathcal{H}) \ln \frac{4em}{\text{Fat}_{\gamma/4}(\mathcal{H})\gamma}}.$$

for some universal constant  $c_0$ .

*Proof.* Theorem 12.8 of Anthony & Bartlett (1999) asserts this for  $m \geq \text{Fat}_{\gamma/4}(\mathcal{H}) \geq 1$  with  $c_0 = 2$ . Now, if  $1 \leq m < \text{Fat}_{\gamma/4}(\mathcal{H})$ , for some universal constant  $c'$ , we have  $\mathcal{N}_\infty(\gamma, \mathcal{H}, m) \leq (c'/\gamma)^m \leq (c'/\gamma)^{\text{Fat}_{\gamma/4}(\mathcal{H})}$ . ■

**Lemma 15 [Theorem 10.1 of Anthony & Bartlett (1999)]** *Suppose that  $\mathcal{H}$  is a set of real-valued functions defined on domain  $X$ . Let  $\mathcal{D}$  be any probability distribution on  $Z = X \times \{0, 1\}$ ,  $0 \leq \epsilon \leq 1$ , real  $\gamma > 0$  and integer  $m \geq 1$ . Then,*

$$\mathbf{P}_{S_m \sim \mathcal{D}} \left[ \exists h \in \mathcal{H} : \text{err}(h) \geq \text{err}_\gamma(h, S_m) + \epsilon \right] \leq 2\mathcal{N}_\infty\left(\frac{\gamma}{2}, \mathcal{H}, 2m\right) e^{-\epsilon^2 m/8},$$

where  $S_m$  is an i.i.d. sample of size  $m$  from  $\mathcal{D}$ .

#### A.4 Proof of Lemma 4

For any fixed  $0 < \gamma < 1/8$  and the given bounded class of distributions with bound  $B \geq 1$ , consider a  $(1/B)$ -bi-Lipschitz base hypothesis class  $\mathcal{H}$  that maps hypothesis from the domain  $X$  to  $[1/2 - 4\gamma, 1/2 + 4\gamma]$ , and define

$$\mathcal{F} := \{x \mapsto h(Mx) : M \in \mathcal{M}, h \in \mathcal{H}\}.$$

Note that finding  $M$  that minimizes  $\text{err}_{\text{hypoth}}$  is equivalent to finding  $f$  that minimizes error on  $\mathcal{F}$ . Using Lemma 19, we have for any  $0 < \gamma < 1/2$ , the sample complexity of  $\mathcal{F}$  is (for all  $0 < \epsilon, \delta < 1/64$ )

$$m \geq \frac{\text{Fat}_{2\gamma}(\pi_{4\gamma}(\mathcal{F}))}{320\epsilon^2}, \quad (10)$$

where  $\pi_{4\gamma}(\mathcal{F})$  is the  $(4\gamma)$ -squashed function class of  $\mathcal{F}$  (see Definition 2 below). We lower bound  $\text{Fat}_{2\gamma}(\pi_{4\gamma}(\mathcal{F}))$  in terms of fat-shattering dimension of  $\mathcal{H}$  to yield the lemma.

To this end we shall first define the  $(\gamma, m)$ -covering and packing number of a generic real-valued hypothesis class  $\mathcal{G}$ . For any domain  $X$ , real-valued hypothesis class  $\mathcal{G} \subset [0, 1]^X$ , margin  $\gamma > 0$ , and a sample  $S \subset X$ , define

$$\begin{aligned} \text{cov}_\gamma(\mathcal{G}, S) &:= \left\{ C \subset \mathcal{G} \mid \begin{array}{l} \forall g \in \mathcal{G}, \exists g' \in C, \\ \max_{s \in S} |g(s) - g'(s)| \leq \gamma \end{array} \right\}, \\ \text{pak}_\gamma(\mathcal{G}, S) &:= \left\{ P \subset \mathcal{G} \mid \begin{array}{l} \forall g \neq g' \in P, \\ \max_{s \in S} |g(s) - g'(s)| \geq \gamma \end{array} \right\} \end{aligned}$$

as the set of  $\gamma$ -covers (resp.  $\gamma$ -packings) of  $S$  by  $\mathcal{G}$ . Let  $\gamma$ -covering number (resp.  $\gamma$ -packing number) of  $\mathcal{G}$  for any integer  $m > 0$  be defined as

$$\begin{aligned} \mathcal{N}_\infty(\gamma, \mathcal{G}, m) &:= \max_{S \subset X: |S|=m} \min_{C \in \text{cov}_\gamma(\mathcal{G}, S)} |C|, \\ \mathcal{P}_\infty(\gamma, \mathcal{G}, m) &:= \max_{S \subset X: |S|=m} \max_{P \in \text{pak}_\gamma(\mathcal{G}, S)} |P| \end{aligned}$$

with the minimizing cover  $C$  (resp. maximizing packing  $P$ ) called as the minimizing  $(\gamma, m)$ -cover (resp. maximizing  $(\gamma, m)$ -packing) of  $\mathcal{G}$ .

With these definitions, we have the following (for some universal constant  $c_0$ ).

$$\begin{aligned} c_0 \left( \frac{m}{16\gamma^2} \right)^{\text{Fat}_{2\gamma}(\pi_{4\gamma}(\mathcal{F})) \ln(em/2\gamma)} &\geq \mathcal{N}_\infty(8\gamma, \pi_{4\gamma}(\mathcal{F}), m) && \text{[Lemma 14]} \\ &\geq \mathcal{P}_\infty(16\gamma, \pi_{4\gamma}(\mathcal{F}), m) && \text{[Lemma 17]} \\ &\geq \left( \frac{1}{32\gamma} \right)^{D^2} \mathcal{P}_\infty(48\gamma, \pi_{4\gamma}(\mathcal{H}), m) && \text{[see (*) below]} \\ &= \left( \frac{1}{32\gamma} \right)^{D^2} \mathcal{P}_\infty(48\gamma, \mathcal{H}, m) && \text{[by the choice of } \mathcal{H}] \\ &\geq \left( \frac{1}{32\gamma} \right)^{D^2} \mathcal{N}_\infty(48\gamma, \mathcal{H}, m) && \text{[Lemma 17]} \end{aligned}$$

$$\geq \left(\frac{1}{32\gamma}\right)^{D^2} e^{\text{Fat}_{768\gamma}(\mathcal{H})/8}. \quad [\text{Lemma 18}] \quad (11)$$

(\*) We show that  $\mathcal{P}_\infty(16\gamma, \pi_{4\gamma}(\mathcal{F}), m) \geq (1/32\gamma)^{D^2} \mathcal{P}_\infty(48\gamma, \pi_{4\gamma}(\mathcal{H}), m)$ , by exhibiting a set  $\mathcal{S} \subset \pi_{4\gamma}(\mathcal{F})$  of size  $(1/32\gamma)^{D^2} \mathcal{P}_\infty(48\gamma, \pi_{4\gamma}(\mathcal{H}), m)$  that is a  $(16\gamma)$ -packing of  $\pi_{4\gamma}(\mathcal{F})$ .

Let  $\pi_{4\gamma}(\mathcal{H}_{48\gamma}) \subset \pi_{4\gamma}(\mathcal{H})$  be a maximal  $(32\gamma)$ -packing of  $\pi_{4\gamma}(\mathcal{H})$  (that is, a maximal set such that for all distinct  $(\pi_{4\gamma} \circ h), (\pi_{4\gamma} \circ h') \in \pi_{4\gamma}(\mathcal{H}_{48\gamma})$ , exists  $x \in X$  such that  $|\pi_{4\gamma}(h(x)) - \pi_{4\gamma}(h'(x))| \geq 48\gamma$ ). Fix  $\epsilon$  (exact value determined later), and define

$$\mathcal{S}_\epsilon := \left\{ x \mapsto (\pi_{4\gamma} \circ h)(Mx) \mid \begin{array}{l} (\pi_{4\gamma} \circ h) \in \pi_{4\gamma}(\mathcal{H}_{48\gamma}), \\ M \in \mathcal{M}_\epsilon \end{array} \right\},$$

where  $\mathcal{M}_\epsilon$  is a  $\epsilon$ -spectral net of  $\mathcal{M}$ , that is, for all  $M \in \mathcal{M}$ , exists  $M' \in \mathcal{M}_\epsilon$  such that  $\sigma_{\max}(M - M') \leq \epsilon$ , and for all distinct  $M', M'' \in \mathcal{M}_\epsilon$ ,  $\sigma_{\max}(M' - M'') \geq \epsilon/2$ .

Then for any two distinct  $f, f' \in \mathcal{S}_\epsilon$ , such that  $f(x) = (\pi_{4\gamma} \circ h)(Mx)$  and  $f'(x) = (\pi_{4\gamma} \circ h')(M'x)$ , we have

- (case 1)  $h$  and  $h'$  are distinct. In this case, there exists  $x \in X$ , s.t.

$$\begin{aligned} |f(x) - f'(x)| &= |\pi_{4\gamma}(h(Mx)) - \pi_{4\gamma}(h'(M'x))| \\ &\geq |\pi_{4\gamma}(h(Mx)) - \pi_{4\gamma}(h'(Mx))| \\ &\quad - |\pi_{4\gamma}(h'(Mx)) - \pi_{4\gamma}(h'(M'x))| \\ &\geq 48\gamma - (1/B)\sigma_{\max}(M - M')\|x\| \\ &\geq 48\gamma - (1/B)\epsilon B = 48\gamma - \epsilon. \end{aligned}$$

- (case 2)  $h, h'$  same but  $M$  and  $M'$  distinct. In this case, there exists  $x$  (with  $\|x\| = 1$ ) s.t.

$$\begin{aligned} |f(x) - f'(x)| &= |\pi_{4\gamma}(h(Mx)) - \pi_{4\gamma}(h(M'x))| \\ &= |h(Mx) - h(M'x)| \\ &\geq B\|(M - M')x\| \\ &\geq B \cdot \min_{M \neq M' \in \mathcal{M}_\epsilon} \sigma_{\max}(M - M') \\ &\geq B(\epsilon/2). \end{aligned}$$

Thus, by setting  $\epsilon = 32\gamma$ , distinct classifiers  $f, f' \in \mathcal{S}_{32\gamma}$  are at least  $16\gamma$  apart (since  $B \geq 1$ ). Hence  $\mathcal{S}_{32\gamma}$  forms a  $(16\gamma)$ -packing of  $\pi_{4\gamma}(\mathcal{F})$ . Therefore, the packing number

$$\mathcal{P}_\infty(16\gamma, \pi_{4\gamma}(\mathcal{F}), m) \geq |\mathcal{S}_{32\gamma}| = |\mathcal{M}_{32\gamma}| |\mathcal{H}_{48\gamma}| \geq (1/32\gamma)^{D^2} \mathcal{P}_\infty(48\gamma, \pi_{4\gamma}(\mathcal{H}), m).$$

Thus, from Eq. (11), it follows that

$$\text{Fat}_{2\gamma}(\pi_{4\gamma}(\mathcal{F})) \geq \Omega\left(\frac{D^2 \ln(1/\gamma) + \text{Fat}_{768\gamma}(\mathcal{H})}{\ln(m/\gamma^2) \ln(m/\gamma)}\right).$$

Combining this with Eq. (10), the lemma follows. ■

**Lemma 16** [ $\epsilon$ -spectral packings of  $D \times D$  matrices] *Let  $\mathcal{M} := \{M \mid M \in \mathbb{R}^{D \times D}, \sigma_{\max}(M) = 1\}$  be the set of matrices with unit spectral norm. Define  $\mathcal{M}_\epsilon \subset \mathcal{M}$  as the  $\epsilon$ -packing of  $\mathcal{M}$ , that is, for every distinct  $M, M' \in \mathcal{M}_\epsilon$ ,  $\sigma_{\max}(M - M') \geq \epsilon$ . Then for all  $\epsilon > 0$ , there exists  $\mathcal{M}_\epsilon$  such that  $|\mathcal{M}_\epsilon| \geq \left(\frac{1}{2\epsilon}\right)^{D^2}$ .*

*Proof.* Fix any  $\epsilon > 0$  and let  $\mathcal{P}_\epsilon$  be a maximal size  $\epsilon$ -packing of Euclidean unit ball  $\mathbf{B}_D$  in  $\mathbb{R}^D$ . That is, for all distinct  $v, v' \in \mathbf{B}_D$ ,  $\|v - v'\| \geq \epsilon$ . Using standard volume arguments (see e.g. proof of Lemma 5.2 of Vershynin (2010)), we know that  $|\mathcal{P}_\epsilon| \geq (\frac{1}{2\epsilon})^D$ . Define

$$\mathcal{M}_\epsilon := \left\{ M' \mid M' = [v'_1 \cdots v'_D] \in \mathbb{R}^{D \times D}, v'_i \in \mathcal{P}_\epsilon \right\}.$$

Then  $\mathcal{M}_\epsilon$  constitutes as an  $\epsilon$ -packing of  $\mathcal{M}$ , since for any distinct  $M, M' \in \mathcal{M}_\epsilon$  such that  $M = [v_1 \cdots v_D]$  and  $M' = [v'_1 \cdots v'_D]$ , we have

$$\sigma_{\max}(M - M') \geq \max_i \|v_i - v'_i\| \geq \epsilon.$$

Without loss of generality we can assume that each  $M \in \mathcal{M}_\epsilon$ ,  $\sigma_{\max}(M) = 1$ . Moreover, by construction,  $|\mathcal{M}_\epsilon| \geq (\frac{1}{2\epsilon})^{D^2}$ . ■

**Lemma 17 [follows from Theorem 12.1 of Anthony & Bartlett (1999)]** *For any real valued hypothesis class  $\mathcal{H}$  into  $[0, 1]$ , all  $m \geq 1$ , and  $0 < \gamma < 1/2$ ,*

$$\mathcal{P}_\infty(2\gamma, \mathcal{H}, m) \leq \mathcal{N}_\infty(\gamma, \mathcal{H}, m) \leq \mathcal{P}_\infty(\gamma, \mathcal{H}, m).$$

**Lemma 18 [Theorem 12.10 of Anthony & Bartlett (1999)]** *Let  $\mathcal{H}$  be a set of real functions from a domain  $X$  to the interval  $[0, 1]$ . Let  $\gamma > 0$ . Then for  $m \geq \text{Fat}_{16\gamma}(\mathcal{H})$ ,*

$$\mathcal{N}_\infty(\gamma, \mathcal{H}, m) \geq e^{\text{Fat}_{16\gamma}(\mathcal{H})/8}.$$

**Lemma 19 [Theorem 13.5 of Anthony & Bartlett (1999)]** *Suppose that  $\mathcal{H}$  is a set of real-valued functions mapping into the interval  $[0, 1]$  that is closed under addition of constants, that is,*

$$h \in \mathcal{H} \implies h' \in \mathcal{H}, \text{ where } h' : x \mapsto h(x) + c \text{ for all } c.$$

*Pick any  $0 < \gamma < 1/2$ . Then for any metric learning algorithm  $\mathcal{A}$  for all  $0 < \epsilon, \delta < 1/64$ , there exists a distribution  $\mathcal{D}$  such that if  $m \leq \frac{d}{320\epsilon^2}$ , then*

$$\mathbf{P}_{S_m \sim \mathcal{D}}[\text{err}(h^*, \mathcal{D}) > \text{err}_\gamma(\mathcal{A}(S_m), \mathcal{D}) + \epsilon] > \delta$$

*where  $d := \text{Fat}_{2\gamma}(\pi_{4\gamma}(\mathcal{H})) \geq 1$  is the fat-shattering dimension of  $\pi_{4\gamma}(\mathcal{H})$ —the  $(4\gamma)$ -squashed function class of  $\mathcal{H}$ , see Definition 2 below—at margin  $2\gamma$ .*

**Definition 2 [squashing function]** *For any  $0 < \gamma < 1/2$ , define the squashing function  $\pi_\gamma : \mathbb{R} \rightarrow [1/2 - \gamma, 1/2 + \gamma]$  as*

$$\pi_\gamma(\alpha) = \begin{cases} 1/2 + \gamma & \text{if } \alpha \geq 1/2 + \gamma \\ 1/2 - \gamma & \text{if } \alpha \leq 1/2 - \gamma \\ \alpha & \text{otherwise} \end{cases}.$$

*Moreover, for a collection  $F$  of functions into  $\mathbb{R}$ , define  $\pi_\gamma(F) := \{\pi_\gamma \circ f \mid f \in F\}$ .*

## A.5 Proof of Lemma 5

Let  $\mathcal{P}$  be the probability measure induced by the random variable  $(\mathbf{X}, Y)$ , where  $\mathbf{X} := (x, x')$ ,  $Y := \mathbf{1}[y = y']$ , st.  $((x, y), (x', y')) \sim (\mathcal{D} \times \mathcal{D})$ .



Define function class

$$\mathcal{F} := \left\{ f_M : \mathbf{X} \mapsto \|M(x - x')\|^2 \mid \begin{array}{l} M \in \mathcal{M} \\ \mathbf{X} = (x, x') \in (X \times X) \end{array} \right\},$$

Following the steps of proof of Lemma 1, we can conclude that the Rademacher complexity of  $\mathcal{F}$  is bounded. In particular,

$$\mathcal{R}_m(\mathcal{F}) \leq 4B^2 \sqrt{\frac{\sup_{M \in \mathcal{M}} \|M^\top M\|_F^2}{m}}.$$

The result follows by noting that  $\phi$  is  $\lambda$ -Lipschitz in the first argument and by applying Lemma 8. ■

## A.6 Proof of Lemma 6

Consider the function class

$$\mathcal{F} := \left\{ f_{v,M} : x \mapsto v \cdot Mx \mid \|v\|_1 \leq 1, M \in \mathcal{M} \right\},$$

and define the composition class

$$\mathcal{F}_\sigma := \left\{ x \mapsto \sum_{i=1}^K w_i \sigma^\gamma(f_i(x)) \mid \begin{array}{l} \|w_i\|_1 \leq 1, \\ f_1, \dots, f_K \in \mathcal{F} \end{array} \right\}.$$

Then, first note that the Gaussian complexity of  $\mathcal{F}$  (with respect to the distribution  $\mathcal{D}$ ) is bounded, since (let  $g_1, \dots, g_m$  denote independent standard Gaussian random variables)

$$\begin{aligned} \mathcal{G}_m(\mathcal{F}, \mathcal{D}) &:= \mathbb{E}_{\substack{x_i \sim \mathcal{D} \\ g_i, i \in [m]}} \left[ \sup_{f_{v,M} \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m g_i f_{v,M}(x_i) \right] \\ &= \frac{1}{m} \mathbb{E}_{\substack{x_i \sim \mathcal{D} \\ g_i, i \in [m]}} \left[ \sup_{\substack{M \in \mathcal{M} \\ \|v\|_1 \leq 1}} v \cdot \sum_{i=1}^m g_i (Mx_i) \right] \\ &= \frac{1}{m} \mathbb{E}_{\substack{x_i \sim \mathcal{D} \\ g_i, i \in [m]}} \left[ \max_j \sup_{M \in \mathcal{M}} \sum_{i=1}^m g_i (Mx_i)_j \right] \\ &\leq \frac{1}{m} \mathbb{E}_{\substack{x_i \sim \mathcal{D} \\ g_i, i \in [m]}} \max_{j \in [D]} \left[ \sum_{i=1}^m g_i \sup_{M \in \mathcal{M}} |(Mx_i)_j| \right] \\ &\leq \frac{c \ln^{\frac{1}{2}}(D)}{m} \mathbb{E}_{x_i \sim \mathcal{D}_X} \max_{j, j' \in [D]} \left( \mathbb{E}_{g_i} \left[ \sum_{i=1}^m g_i \left( \sup_{M \in \mathcal{M}} |(Mx_i)_j| - \sup_{M' \in \mathcal{M}} |(M'x_i)_{j'}| \right) \right]^2 \right)^{\frac{1}{2}} \\ &= \frac{c \ln^{\frac{1}{2}}(D)}{m} \mathbb{E}_{x_i \sim \mathcal{D}_X} \max_{j, j' \in [D]} \left( \sum_{i=1}^m \left[ \sup_{M \in \mathcal{M}} |(Mx_i)_j| - \sup_{M' \in \mathcal{M}} |(M'x_i)_{j'}| \right]^2 \right)^{\frac{1}{2}} \\ &\leq c' B \sqrt{\frac{d \ln D}{m}}, \end{aligned}$$

where (i) second to last inequality is by applying Lemma 20, (ii)  $c, c'$  are absolute constants, (iii)  $d := \sup_{M \in \mathcal{M}} \|M^\top M\|_F^2$ . Note that bounding the Gaussian complexity also bounds the Rademacher complexity by Lemma 21.

Finally by noting that  $\mathcal{F}_\sigma$  is a  $\gamma$ -Lipschitz composition class of  $\mathcal{F}$  and  $\phi^\lambda$  is a classification based loss function that is  $\lambda$ -Lipschitz in the first argument, we can apply Lemma 8 yielding the desired result. ■

**Lemma 20 [Lemma 20 of Bartlett & Mendelson (2002)]** *Let  $Z_1, \dots, Z_D$  be random variables such that each  $Z_j = \sum_{i=1}^m a_{ij}g_i$ , where each  $g_i$  is independent  $N(0, 1)$  random variables. Then there is an absolute constant  $c$  such that*

$$\mathbb{E}_{g_i} \max_j Z_j \leq c \ln^{\frac{1}{2}}(D) \max_{j,j'} \sqrt{\mathbb{E}_{g_i} (Z_j - Z_{j'})^2}.$$

**Lemma 21 [Lemma 4 of Bartlett & Mendelson (2002)]** *There are absolute constants  $c$  and  $C$  such that for every class  $\mathcal{F}$  and every integer  $m$*

$$c\mathcal{R}_m(\mathcal{F}, \mathcal{D}) \leq \mathcal{G}_m(\mathcal{F}, \mathcal{D}) \leq C \ln(m) \mathcal{R}_m(\mathcal{F}, \mathcal{D}),$$

where  $\mathcal{R}$  and  $\mathcal{G}$  are Rademacher and Gaussian complexities of a function class  $\mathcal{F}$  with respect to the distribution  $\mathcal{D}$  respectively.

## A.7 Proof of Corollary 7

The conclusion of Eq. (3) is immediate by dividing the given failure probability  $\delta$  across the sequence  $\mathcal{M}^1, \mathcal{M}^2, \dots$  such that  $\delta \mu_d$  failure probability is associated with class  $\mathcal{M}^d$ , then apply Lemma 5 (for distance based metric learning) or Lemma 6 (for classifier based metric learning) to each class  $\mathcal{M}^d$  individually, and finally combining the individual deviations together with a union bound.

For the second part, for any  $M \in \mathcal{M}$  define  $d_M$  and  $\Lambda_M$  as per the lemma statement. Then with probability at least  $1 - \delta$

$$\begin{aligned} \text{err}^\lambda(M_m^{\text{reg}}, \mathcal{D}) - \text{err}^\lambda(M^*, \mathcal{D}) &\leq \text{err}^\lambda(M_m^{\text{reg}}, S_m) + d_{M_m^{\text{reg}}} \Lambda_{M_m^{\text{reg}}} - \text{err}^\lambda(M^*, \mathcal{D}) \\ &\leq \text{err}^\lambda(M^*, S_m) + d_{M^*} \Lambda_{M^*} - \text{err}^\lambda(M^*, \mathcal{D}) \\ &\leq O(d_{M^*} \Lambda_{M^*}) = O(\epsilon), \end{aligned}$$

where (i) the first inequality is by applying Eq. (3) on weighting metric  $M_m^{\text{reg}}$  (with failure probability set to  $\delta/2$ ), (ii) the second inequality is by noting that  $M_m^{\text{reg}}$  is the (regularized) sample error minimizer as per the lemma statement, (iii) the third inequality is by applying Eq. (3) on weighting metric  $M^*$  (with failure probability set to  $\delta/2$ ), and (iv) the last equality by noting the definitions of  $\Lambda_{M^*}$  and our choice of  $m$ . ■